Copy 45 of 72 copies

AD-A225 518

IDA DOCUMENT D-780

PANEL REVIEW OF LONG-HAUL NETWORKING IN DISTRIBUTED SIMULATION

Jeffrey D. Case
Danny Cohen
Dale B. Henderson
Irwin L. Lebow
David L. Mills
J. D. Fletcher

June 1990

Prepared for
Defense Advanced Research Projects Agency

INSTITUTE FOR DEFENSE ANALYSES
1801 N. Beauregard Street, Alexandria, Virginia 22311-1772

IDA Log No. HQ 90-35561

DEFINITIONS

IDA publishes the following documents to report the results of its work.

Reports

Reports are the most authoritative and most carefully considered products IDA publishes. They normally embody results of major projects which (a) have a direct bearing on decisions affecting major programs, (b) address issues of significant concern to the Executive Branch, the Congress and/or the public, or (c) address issues that have significant economic implications. IDA Reports are reviewed by outside panels of experts to ensure their high quality and relevance to the problems studied, and they are released by the President of IDA.

Group Reports

Group Reports record the findings and results of IDA established working groups and panels composed of senior individuals addressing major issues which otherwise would be the subject of an IDA Report. IDA Group Reports are reviewed by the senior individuals responsible for the project and others as selected by IDA to ensure their high quality and relevance to the problems studied, and are released by the President of IDA.

Papers

Papers, also authoritative and carefully considered products of IDA, address studies that are narrower in scope than those covered in Reports. IDA Papers are reviewed to ensure that they meet the high standards expected of refereed papers in professional journals or formal Agency reports.

Documents

IDA Documents are used for the convenience of the sponsors or the analysts (a) to record substantive work done in quick reaction studies, (b) to record the proceedings of conferences and meetings, (c) to make available preliminary and tentative results of analyses, (d) to record data developed in the course of an investigation, or (e) to forward information that is essentially unanalyzed and unevaluated. The review of IDA Documents is suited to their content and intended use.

The work reported in this document was conducted under contract MDA 903 89 C 0303 for the Department of Defense. The publication of this IDA document does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that Agency.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
June1990

3. REPORT TYPE AND DATES COVERED
Final--February to March 1990

4. TITLE AND SUBTITLE
Panel Review of Long-Haul Networking in Distributed Simulation

5. FUNDING NUMBERS
C - MDA 903 89 C 0003
T - A-132

6. AUTHOR(S)
Jeffrey D. Case, Danny Cohen, Dale B. Henderson, Irwin L. Lebow, David L. Mills, J.D. Fletcher

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA 22311-1772

8. PERFORMING ORGANIZATION REPORT NUMBER
IDA Document D-780

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Defense Advanced Research Projects Agency
Material Sciences Division
1400 Wilson Boulevard
Arlington, VA 22209

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The report documents a review of the technical approach used to support long-haul, networked data communications in DARPA's Advanced Distributed Simulation Technology (ADST) Program. The review was conducted 1-2 March 1990 by an independent panel of five scientists whose comments are presented and summarized. The panel concluded that: (1) the current technical approach to long-haul data communication in ADST is sound, given current system requirements and resources; (2) the existing architecture and protocols will support a one order of magnitude increase in traffic over the next five years, assuming the current pace of improvements in computation continues and is reflected in system improvements; (3) the ability of the current approach to support a two order of magnitude increase in traffic over the next 10 years is much less certain; (4) the system has been appropriately designed to support migration to standardized architectures and protocols; (5) the system should incorporate standardized approaches where they satisfy the levels of performance needed, but performance should receive preference over standardization where they do not; (6) a comprehensive assessment of system requirements for long-haul data communications requirements in the network is needed and should be undertaken promptly; (7) general purpose mechanisms for network management and security should be incorporated in the system; (8) data collection and analysis of network traffic and requirements should be instituted as a routine component of system operations to monitor performance and to guide system modifications and growth.

14. SUBJECT TERMS
Data Communications, Engagement Training, Networking, SIMNET, Simulation, Simulator Networking, Training Simulation

15. NUMBER OF PAGES
139

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	SAR

IDA DOCUMENT D-780

PANEL REVIEW OF LONG-HAUL NETWORKING IN DISTRIBUTED SIMULATION

Jeffrey D. Case
University of Tennessee

Danny Cohen
U.S.C./Information Sciences Institute

Dale B. Henderson
Strategic Defense Initiative Organization
National Test Bed

Irwin L. Lebow
Private Consultant

David L. Mills
University of Delaware

J. D. Fletcher
Institute for Defense Analyses

June 1990

INSTITUTE FOR DEFENSE ANALYSES

Contract MDA 903 89 C 0003
DARPA Assignment A-132

ABSTRACT

This report documents a review of the technical approach used to support long-haul, networked data communications in DARPA's Advanced Distributed Simulation Technology (ADST) Program. The review was conducted 1-2 March 1990 by an independent panel of five scientists whose comments are presented and summarized. The panel concluded that: (1) the current technical approach to long-haul data communication in ADST is sound, given current system requirements and resources; (2) the existing architecture and protocols will support a one order of magnitude increase in traffic over the next five years, assuming the current pace of improvements in computation continues and is reflected in system improvements; (3) the ability of the current approach to support a two order of magnitude increase in traffic over the next 10 years is much less certain; (4) the system has been appropriately designed to support migration to standardized architectures and protocols; (5) the system should incorporate standardized approaches where they satisfy the levels of performance needed, but performance should receive preference over standardization where they do not; (6) a comprehensive assessment of system requirements for long-haul data communications requirements in the network is needed and should be undertaken promptly; (7) general purpose mechanisms for network management and security should be incorporated in the system; (8) data collection and analysis of network traffic and requirements should be instituted as a routine component of system operations to monitor performance and to guide system modifications and growth.

CONTENTS

ABBREVIATIONS

ACSE	Association Control Service Element
ACTS	Advanced Communications Technology Satellite
ADST	Advanced Distributed Simulation Technology
ASN.1	Abstract Syntax Notation One
BBN	Bolt Beranek and Newman
CCTT	Close Combat Tactical Trainer
CLNP	Connection-Less Network Protocol
CPU	Central Processing Unit
DARPA	Defense Advanced Research Projects Agency
DAX	Digital Automatic Cross-Connect
DETEC	Defensive Technology Evaluation Code
DoD	Department of Defense
ECM	Electronic Countermeasures
ECCM	Electronic Counter-Countermeasures
FAF	Fully Automated Forces
FOV	Field of View
I/O	Input/Output
IP	Internet Protocol
ISO	International Standards Organization
LAN	Local Area Network
LHN	Long-Haul Network
NASA	National Aeronautics and Space Administration

NSF	National Science Foundation
NTB	National Test Bed
OSI	Open Systems Interconnection
OSPF	Open-System Shortest Path First
PDU	Protocol Data Unit
POV	Point of View
PPS	Packets Per Second
R&D	Research and Development
RISC	Reduced Instruction Set Computer
SAF	Semi-Automated Forces
SDI	Strategic Defense Initiative
SIMNET	Simulator Network
SNMP	Simple Network Management Protocol
ST	Stream Transport
VSAT	Very Small Aperture Terminal
VTC	Video Terminal Connections
WAN	Wide Area Network

SUMMARY

A. INTRODUCTION

At the request of the Defense Advanced Research Projects Agency (DARPA), the Institute for Defense Analyses (IDA) conducted a peer review of the technical approach being used to support long-haul, networked data communications in the Advanced Distributed Simulation Technology (ADST) Program.

Members of the review panel were the following:

Jeffrey D. Case
Associate Professor of Computer Science
University of Tennessee

Danny Cohen
Director, Systems Division
USC/Information Sciences Institute

Dale B. Henderson
Chief Scientist,
Strategic Defense Initiative Organization National Test Bed

Irwin L. Lebow
Private Consultant

David L. Mills
Professor of Electrical Engineering
University of Delaware

Brief biographical sketches of the panel members are provided in Appendix B.

The panel was asked to consider the following question:

With regard to the development and implementation of long-haul networking in the Advanced Distributed Simulation Technology Program is there anything, within obvious constraints of time, budget, and available resources, that should be done to better meet the goals of the program?

About two weeks prior to the review, each panel member received a "read-ahead" package that provided information on DARPA's ADST program and the approaches used to support its long-haul, networked data communications. The documents included in this package are listed in Appendix C.

The review was held on 1-2 March 1990 at the ADST Facility in Rosslyn, Virginia. The agenda for the review covered four basic activities:

- Orientation, statement of panel responsibilities, description of ADST program objectives and technologies.

  [LTC (P) James Shiflett (DARPA) and program staff.]

- Description and discussion of ADST long-haul network objectives, development, and technical approaches.

  [Duncan Miller, Steven Blumenthal, Jerry Burchfiel, Alexander MacKenzie, Arthur Pope, Claudio Topolcic, and Rolland Waters, all of Bolt Beranek and Newman (BBN)].

  The briefing materials used in the BBN presentations are included as Appendix E.

- Discussion within the panel and review of the technical approaches taken to support long-haul networking in the ADST program. These discussions were preceded by an hour-long briefing to the panel by Michael Sabo of SSDS, Inc., representing the Martin Marietta Team for Distributed Simulators Architecture. The briefing materials used in this presentation are included as Appendix F.

- Informal report by the panel to DARPA and other DoD representatives on preliminary findings concerning ADST long-haul networking approaches.

Those who attended the review meetings are listed in Appendix D.

B. FINDINGS

After the review, members of the panel documented their comments and recommendations. A summary of these comments follows.

1. Quality of the Technical Approach

Three of the panelists discussed the overall quality of the technical approach that has been taken and the work completed thus far. All three found the current architecture for long-haul networking to be sound and the necessary design trade-offs to be appropriate given the objectives and constraints of the program.

2. Adequacy of the Approach for Projected Growth

All five panelists discussed the capacity of the current architecture and protocols to support foreseeable growth in DoD use of advanced distributed simulation. The following notional growth function was used to assess long-haul communications requirements and to concretize the discussion:

Year	Sites	Objects
1990	5	1,000
1992	50	30,000
1994	100	10,000
2000	300	100,000

This function and its rough form--one order of magnitude in five years, two orders of magnitude in 10--were discussed. Generally, the panel concluded that a growth of one order of magnitude in about five years would be supported by the current architecture, assuming that the current pace of development and improvements in computer and communication technology continues and that the rate with which these improvements are incorporated into the distributed simulation system also continues.

The views of the panel regarding a growth of two orders of magnitude in 10 years were mixed, reflecting different assumptions concerning modifications and improvements in the technologies and algorithms used. The panel concluded that two orders of magnitude of growth in 10 years is a reasonable and attainable goal, but that substantial changes in the current architecture and algorithms are needed to support it.

3. Performance versus Standardization

Four of the panelists discussed this issue. All four noted the value of standardization in general. Two of the four recommended that DARPA encourage the development of standard interfaces for communicating information concerning vehicle dynamics and visual displays.

However, all four panelists emphasized the need of the distributed simulation system to satisfy current performance criteria and the consequent necessity of incorporating application specific approaches that satisfy these criteria in place of standardized approaches that do not. The panelists noted that BBN has attempted to establish and maintain evolutionary paths to emerging standard approaches that may satisfy the performance criteria of the system. These efforts appear to be sound as exemplified by the ease with which DoD Internet Stream Transport (ST) could be replaced in the current approach.

Three panelists emphasized that standardization is more important to external interfaces, such as those required by the network services, than to internal interfaces peculiar to distributed simulation. They suggested that the system protocol data units (PDUs) should *not* be required to comply with the encoding rules of Abstract Syntax Notation One (ASN.1).

4. LHN Needs Assessment

All the panelists recommended that DARPA soon complete a systematic review and assessment of capabilities needed to support long-haul networking for distributed simulation. Two panelists recommended that existing programs be combed for emerging and applicable techniques for meeting needs identified by the assessment. Two panelists registered concern that satellite communications may prove to be impracticable for distributed simulation -- or at least for this application as it is currently implemented. Suggestions for meeting the needs of long-haul networking in ADST included the following:

- Generate objects simulated for semi-automated and fully-automated forces locally;
- Compress PDUs by tailoring them more closely to their application;
- Aggregate PDUs;
- Filter PDUs at gateways through redesigned multicast addresses, hierarchical structuring with intermediate processors, or otherwise;
- Devise more sophisticated dead reckoning algorithms for extrapolating vehicle movements;
- Provide association management for subnets;
- Tailor data communication reliability to specific performance requirements;
- Examine alternatives to ST, including Internet Protocol (IP) multicasting with and without special purpose, resource reservation services built into the network;
- Encode PDUs directly into IP datagrams with IP multicast addresses.

5. Management and Security

Four panelists noted the inadequacy/absence of mechanisms for security. Currently there is no mechanism for a simulator to determine the authenticity of an arriving PDU, and the mechanisms for preventing conflict between two independent simulations running simultaneously are weak. The panel recommended that these problems be remedied through use of general purpose architecture -- Simple Network Management Protocol (SNMP) was specifically mentioned by two panelists as a possibility.

6. Data Collection

Two panelists mentioned the absence of data collection and analysis to characterize system transport requirements and the traffic matrix. They recommended instituting routine procedures for this purpose.

The complete written comments of the panelists are provided in Appendix A.

C. RECOMMENDATIONS

Recommendations based on the the panel's findings include the following:

- DARPA should continue to emphasize performance over standardization in ADST long-haul networking and incorporate standard architectures and protocols when it is possible to do so without compromising performance.
- A systematic and comprehensive needs assessment for long-haul, networked data communications in ADST should be conducted soon. System architecture and protocols should be modified based on this assessment.
- A general purpose, standard mechanism for network management should be incorporated in the system.
- A general purpose, standard mechanism for network security should be incorporated in the system.
- Data collection and analysis of network functioning should be included as a routine component of system operations.

APPENDIX A
PANEL MEMBERS' COMMENTS

COMMENTS ON LONG-HAUL NETWORKING IN ADVANCED DISTRIBUTED SIMULATION TECHNOLOGY

Jeffrey D. Case
University of Tennessee

The consensus of the group is that the SIMNET system is basically sound, for it obviously works in the environments where it has been exercised. For this to be true, the networking component of SIMNET must also be sound and the evidence to date indicates that soundness.

There are several issues which were explored in some depth (to the extent time permitted):

A. ARCHITECTURE

The long-haul network component of the SIMNET system appears to have been added late in the development cycle and perhaps might have been architected into design differently had its need been envisaged earlier. However, the long-haul component of the system does appear to operate in networks of the current scale and its operational readiness speaks loudly as a genuine gauge of its worth.

However, there should be a thorough review of the basic SIMNET model with a careful examination of the basic simulation model to determine if the model which was originally selected is still appropriate now that long-haul networking is an important component. It is reasonable to believe that the fact that there is a long haul network should have impacts on the simulation architecture.

Two important areas should be given special attention.

First, consideration should be given to generating the simulated objects in a distributed fashion. This would entail broadcasting the commands to generate the simulated objects and then generating them on replicated systems at each local simulation site, perhaps only generating those which are potentially of interest at that site rather than broadcasting information about the simulated objects. This would require that the generators of the simulated forces act in a deterministic manner. The panel was led to understand that there is a fundamental problem with the simulation model which makes such a plan unworkable. It may be that minor modifications to the simulation model such as moving the referee function to a third party can relax those restrictions. If they can be relaxed, there are significant opportunities for reducing the bandwidth requirements on the long haul portion of the network. It is important to note that the number of simulated objects is likely to become an increasing percentage of the objects participating in simulations.

Second, the simulation model appears to be sensitive to time delays, and consequently the use of satellite transmission is impractical. This appears to be a serious problem, especially if shipboard simulations are to become a reality.

B. SCALABILITY AND GROWTH

The panel gave major attention to the questions/issues surrounding scalability and growth. We assumed the following projections (as given):

Year	Sites	Objects
1990	5	1,000
1992	50	3,000
1994	100	10,000
2000	300	100,000

The above figures show a growth requirement of approximately one order of magnitude in numbers of packets and systems over the next 2 to 5 years and approximately two orders of magnitude over the next decade.

It is believed that the 2-5 year requirement (one order of magnitude in numbers of packets and systems) can be met. It is believed that the general improvements in technology that will naturally come with time [customary price/performance improvements in processor, memory, and input/output (I/O) technology] will be sufficient to meet the short-term goals.

However, this conclusion is mitigated by other increases in traffic which may be generated by other aspects of the application. It is understood that there are multiple opportunities to extend the simulation to new areas, such as intelligence and logistics, which would potentially have significant impacts on the traffic matrix and thereby render this conclusion inaccurate. However, if these new portions of the application do not materially alter the traffic matrix which is currently dominated by vehicle appearance PDUs, their impact should be slight.

It is believed that the long-term goal which will require two orders of magnitude of growth cannot be achieved without architectural changes such as:

1. Use of application level gateways to provide increased filtering at the boundary between the local area networks and the long-haul network
2. Increased use of filtering and parallelism in the generation of simulated forces.

Research is needed to identify the changes that will be required.

C. STANDARDIZATION

The panel also invested time in examination of the questions/issues surrounding standardization. Standardization is often touted as important because it can lead to:

- Reduced development costs
- Reduced development time / time to market
- Reduced replication costs resulting from the availability of off-the-shelf (commodity) items and competition from second sources.

It is important to note that interoperability and standardization are not the same thing and that these issues are often confused with layering and modularization.

While standardization is often a good thing, it is important to ask for what purpose, and at what cost? Regarding the benefits above, the development costs are largely sunk investment, and since the development is now largely completed, requiring conversion to standards (which in some cases are yet to be defined) would only increase the length of the development cycle. However, the replication costs are likely to become significant in light of the 1 to 2 orders of magnitude of growth that are anticipated.

It is expected that the cost of individual simulators will, due to sheer quantities, dominate the replication costs. Investments to allow replication of these systems to widely available hardware and software platforms based on standards are worthy of investigation.

At first blush, it appears that increased standardization in the long-haul portion of the network is not likely to yield significant benefits or the availability of second sources. In any case, it is likely that the long-haul component will be acquired from a single source.

However, some further thought leads to a different conclusion. The SIMNET simulators are likely to be located at installations which have unrelated investments in network infrastructure. If the simulators are dispersed within a particular site, the network infrastructure within the site will need to be compatible with the SIMNET protocols. Consequently, the use of nonstandard protocols within SIMNET is likely to be problematic within those installations which have a geographic dispersion of the simulators and which also want to use networking equipment to support standard protocols.

It is worthwhile to note that neither the arguments of the presenters who depicted ST and its descendant as popular and widely established standards nor the arguments of the presenters who campaigned for migration to the Open Systems Interconnection/ International Standards Organization (OSI/ISO) model and protocol suite were found to be especially credible or compelling. (In fairness to the presenter on Friday morning, I was unable to attend the first portion of the presentation and I may have misunderstood the points being made).

There are often performance penalties associated with standardization because general-purpose protocols are seldom optimized for a specific purpose and seldom are as efficient. For example, the Abstract Syntax Notation One (ASN.1) provides a general solution for the problems encountered in multivendor networks of heterogeneous systems including word size, byte ordering, arithmetic types, and character set differences. However, the use of Abstract Syntax Notation One (ASN.1) is inappropriate for use with the vehicle appearance PDUs because the performance of existing ASN.1 encoders and decoders is incompatible with the required packet rates.

In an attempt to be clear and unambiguous, the use of ASN.1 for the vehicle appearance PDUs is NOT recommended. It may be worthwhile to use standards such as ASN.1 for other messages such as those for network management and for managing the distributed simulation.

The discussion of these standardization issues is especially timely as the transition from the research environment to operational deployment can be an effective opportunity for their resolution. It is possible to standardize too early or too late.

D. NAGGING CONCERNS

Having said the above, there are additional areas of nagging concerns that warrant identification and some discussion. These are especially present in the areas of security and management.

The system, as presently designed and implemented, is totally devoid of a security architecture. In this context, security is meant to mean more than just privacy.

There appear to be no mechanisms for a simulator to determine the authenticity of an arriving PDU. As a result. the SIMNET network is vulnerable to a number of threats, accidental or intentional. In many ways, the network is similar to the architecture of a large bridged Ethernet with full broadcast capability and, as such, suffers from many of the drawbacks which are inherent in that environment and well understood.

The mechanisms for managing the simulation appear to be especially weak. In particular, it is possible, and even easy, for there to be conflicts between two independent simulations (i.e., it is possible for a station joining the exercise to find itself in the middle of the wrong war, to the detriment of both exercises). Consequently, there is a need for additional development of the necessary protocols, tools, and procedures to manage these aspects including such tasks as address management.

To the maximum extent possible, these development efforts should attempt to follow the models used in general purpose networks with a general purpose security architecture.

It is difficult to overstate the importance of a good network management subsystem in operational networks of the size of SIMNET. One presentation proposed use of the Simple Network Management Protocol (SNMP) for this purpose. The SNMP seems well-suited in this application (but the panelist's biases should be obvious).

E. RESEARCH ISSUES

One question that the review panel was charged with answering is "What research areas should be pursued and supported?" There are several technical refinements which should be pursued in the short term and the long term. Some have already been mentioned.

In the short term, it is important to develop specifications for the long-haul portion for the Army Close Combat Tactical Trainer (CCTT) procurement. Of necessity, the design should mirror that which is already implemented.

In the longer term, work should be undertaken to engineer solutions which will realize the multicasting and guaranteed service objectives in light of the recent progress made on IP multicasting related to Open-System Shortest Path First (OSPF) development. It is hoped that these efforts will allow SIMNET to make greater use of off-the-shelf networking equipment. Alternatively, work on the evolution from ST1 to ST2 must be undertaken. Finally, data collection and analysis to provide characterization of the transport requirements and the traffic matrix should be pursued so that planning for growth can be based on sound data.

These research activities should be performed making use of the applicable experience in the private sector for similar problems. The research will also have spinoff contributions to those problems in the commercial sector.

For example, the data communications challenges associated with the operation of stock and monetary exchanges are analogous to the SIMNET problem. In both cases, widely geographically and organizationally separate personnel need to observe and interact with a shared view of a globally distributed database, issuing transactions in real time or near real time. The problems of network management and operation, including the management of subscriber enrollment/engagement, are quite similar. The SIMNET project can both benefit from and contribute to this growing body of related knowledge in future research activities.

COMMENTS ON LONG-HAUL NETWORKING IN SIMNET

Danny Cohen
Information Sciences Institute

1. The present architecture proved to be very good. It performs a very good job of separating the issues (aka layering). Practically all the simulation-related issues are contained in the upper level, the application layer(s). The ease with which many new capabilities were added to the system, and the ethernet-based architecture which was expanded to operate over DoD Internet Stream Transport (ST) (requiring changes neither to the application nor to ST) proved the architecture. Both the simulation architecture and the communication architecture are to be commended.

ST could be replaced by any other reasonable protocol supporting multicast and capable of guaranteeing the requested performance. This replacement should be relatively easy thanks to the simulation and communication architectures.

2. It is my assessment that this architecture (with possible minor improvements) would be able to support both growth of 10X (in 5 years) and 100X (in 10 years) in the number of sites and objects.

This assessment is based on the assumption that both computing power and communication performance will keep progressing at the same rate as they did in recent years.

In addition, improvements in algorithms and the system architecture will contribute to this goal.

This assessment is based primarily on the simulation as it exists today, without considering additional services that may have to be supported such as multiple VTCs and other capabilities such as electronic countermeasures (ECM) and electronic counter-countermeasures (ECCM).

Other helping assumptions are: (a) The geographic size of exercise grows with the number of objects, such that the geographic density of objects does not increase with the total number of objects; (b) Most inter-object interactions are limited in their geographic size, hence with proper organization of exercises the communication load grows less than linearly with the number of objects; (c) The percentage of fully-automated forces (FAFs) and semi-automated forces (SAFs) will grow as the total number of objects grows. This will reduce further the load on the communication, as discussed in 3(e).

3. I recommend that the implication of long-haul communication be thoroughly examined, and if so desired the system be modified accordingly.

The introduction of long-haul communication has implications both on the communication and on the simulation. For example:

(a) It may be desirable to use some application-specific encoding (compression) of PDUs to decrease their size in order to reduce the load on the long-haul network (LHN).

(b) It may be desirable to aggregate PDUs in order to form fewer larger packets.

(c) It may be desirable to define multicast addresses in such a way that will support some filtering to eliminate the transmission of PDUs to destinations that don't need them. [BBN has started doing that.]

(d) It may be desirable to reduce the frequency of inter-site appearance PDU updates, below the rate used intra-site.

(e) It may be desirable to replicate in all sites the processing of SAFs and FAFs such that only their manual input has to be communicated over LHNs, not their PDUs. This will require some reliable mechanism (e.g., with acknowledgment) for the state-modifying PDUs (such as "being killed"). Luckily, these messages constitute only a very small percentage of the total traffic.

In the above, (a) thru (c) are communication issues, whereas (d) and (e) are simulation issues. The current architecture supports an application-specific "Intelligent Gateway" which could perform the above (a) thru (d).

The implementation of (e) would require some changes to the existing system and maybe also a minor modification to its architecture.

4. It is not a state secret that "standardization" and "efficiency" are not synonymous. It is clear that the system must achieve a significant performance in order to meet its real-time requirements. I dare suggest that any excess capacity will always be used by the system developers for various enhancements and additional features. Hence, performance will always be, by definition (nearly), a critical force driving the development of the system.

Therefore, I submit that the system should be developed to be as efficient as possible, while confining standardization to its external interfaces. This would guarantee its interoperability with other systems without penalizing its performance by insisting that standards be used internally throughout the entire system, especially with standards that are not "combat proven."

To be specific, intra-system PDUs should not be forced to comply with ASN.1 !!!

Similarly, the use of networks should comply with the prevailing standards in the communication systems in use.

5. I strongly urge DARPA to push for a standard interface between the simulators of vehicle dynamics and their visual subsystems.

Two types of information have to be communicated over this interface: viewing of the "world" and modification of the "world." The former defines the viewing parameters (e.g., visibility, POV, FOV, illumination, and display characteristics), and the latter defines the position, orientation, and status of the dynamic objects to be viewed.

Standardization of this interface will increase significantly the vendor base for the visual components of simulators. These components typically require a substantial portion of the total cost of the simulation system.

I expect such a standard interface to have a significant payoff even in the short run by separating the vendors into application-specific vendors (e.g., for the dynamics, weapon systems, and mockups) from the more general imagery vendors. Such a separation will allow the government to procure independently the bests of both worlds, regardless of teaming arrangements.

(At present the large vendors of the visual subsystems don't bother bidding on specific systems that require special development and are not expected to be purchased in large quantities, e.g., for a specific guided missile.)

6. In order to protect the system from potentially over-creative vendors, I'd suggest adopting the requirement that the system use a general purpose communication and general purpose security architecture. The requirements of this system do not justify the development of special purpose communication schemes or special purpose security architecture.

7. In summary, nothing in the present architecture excludes the implementation of any of the suggestions presented above.

It is my assessment that the present architecture (possibly with minor modifications and upgrades) is capable of supporting an order of magnitude growth in the number of site and objects over the next 5 years, and most likely another order of magnitude over the next 5 years (i.e., 100X by the year 2000).

COMMENTS ON LONG-HAUL NETWORKING

Dale Henderson
SDIO National Test Bed

The principal question was whether we foresaw any barrier to the expansion of SIMNET from its present scale to 10 times as many objects. I agree with the consensus that this expansion should be possible with only minor structural changes to the simulation and to the network and with the natural evolution of computing hardware. However, I am surprised that we were not presented with any empirical data on the delays and loading experienced under various numbers of simulation objects.

I would think that with the semi-automatic objects it would be easy and inexpensive to create a great deal of message traffic just by adjusting the threshold for dead reckoning updates and by adjusting the numbers of objects. One could then studv the message delays experienced by a very few discrete objects both as now implemented and with the proposed communications processor and by those located both on the same LAN and over a long-haul connection. With these data one could be able to find the various "knees" of the present system and better predict the performance with 10 or 100 as many objects.

The basic SIMNET paradigm is simple, elegant, and successful; I rather admire it. But, given the great expansion in the numbers and types of simulation objects, I wonder whether a review of the paradigm might not have been in order even without the long-haul question. It may be time to consider a hierarchical structure with intermediate or "referee" processors to filter, respond, or reformat messages. Damage assessment, for example, might be assigned to such an intermediate. So might physics-intensive tasks such as signal propagation through (say) nuclear effects.

The proposal to use the gateway processors for such filtering tasks appears to have merit. We should at least recognize that the tasks are logically separate. And some intermediate filtering may be useful in places besides long-haul gateway nodes.

I commend the project management for including the second-day speaker because he represents a strongly held dissenting position. However, I found his argument to be wholly unconvincing.

Several high-bandwith extra-simulation applications were mentioned as additional ways to justify the cost of expensive communications--video teleconferencing, for example. I doubt this is necessary or wise. It is probably better to use commercially available teleconferencing.

Security--in all of its aspects--was raised by the speakers and by the panel. I agree that SIMNET ought to have a security effort, and think that this could be in collaboration with other projects such as the SDI National Test Bed, possibly producing more for everybody at less expense.

I conclude with a brief comparison of SIMNET and the simulation at the National Test Bed Defense Technical Evaluation Code (DETEC). SIMNET was designed for many-person training with analysis as a subsidiary mission; DETEC was designed with the inverse emphases. DETEC was designed for supercomputers; SIMNET has no central computer. SIMNET's main communication is ethernet; DETEC's is shared memory.

Despite these contrasts, the two simulation philosophies and code structures are quite similar. Both are object oriented in philosophy. Both are careful to confine inter-object communication to explicit messages. It is interesting that both groups independently came to similar designs. This means that interoperation of the two systems should be fairly easy if there were a reason to attempt it.

COMMENTS ON SIMNET LONG-HAUL NETWORKING

Irwin L. Lebow
Consultant, Washington, D.C.

A. INTRODUCTION

The SIMNET program has been under way at DARPA since 1983. It started out by linking collocated simulators and then extended its scope to include remotely located simulators by linking the local clusters to one another with long-haul telecommunications. This year the technology is to be transferred to the Army, with the expectation of a great expansion in the next decade. Because the long-haul communications in the existing R&D program were introduced in an ad hoc way relatively late in the program, DARPA management convened our panel to review the long-haul communications and thereby either assure itself that it was suitable for transfer or, if not, make recommendations for changes.

More specifically, DARPA asked the panel to comment on the suitability of the long-haul approach to sustain SIMNET growth levels of one order of magnitude in five years and two orders of magnitude in 10 years.

The panel addressed DARPA's specific request and, of course, unearthed many associated issues. My comments are largely addressed to the main question on which the panel was in agreement.

B. THE IDEA OF SIMNET

SIMNET is a system that permits training of personnel and evaluation of doctrine in battlefield situations. It is not a technique for training a soldier how to operate a particular vehicle; all of the Services have had such simulators for many years. It is rather a technique for training a vehicle operator to participate in a battle. A system such as SIMNET that permits training at this higher level has great potential for the future with its expected austere fiscal environment.

C. THE SIMNET LONG-HAUL COMMUNICATIONS CONCEPT

The simulators are physically located in geographically dispersed clusters, communicating with one another on local area networks (LANs) interconnected with long-haul transmission. They keep track of one another through the exchange of position report packets. Using these reports, each simulator performs a simple dead-reckoning calculation to predict future positions. This dead-reckoning permits transmission of the packets at a reduced rate and, in addition, permits the communications to be relatively unreliable (no acknowledgements).

It is easy and economical for each simulator to send a periodic report of its position to every other simulator on the same LAN regardless of its simulated location. However, it is necessary to perform a filtering operation to restrict long-haul reporting to those simulators that are within sight of one another using gateway processors at the termini of the interconnecting circuits for this purpose. The position reporting must be multicast for speed and efficiency. This is a relatively simple matter for inherently point-to-multipoint media such as satellites, but complex for inherently point-to-point media such as landlines. Nevertheless SIMNET has used the latter to avoid having to cope with the transmission delay of satellites. It does this using the multicast stream protocol, ST, implemented at the gateways.

D. ANALYSIS OF DARPA'S REQUEST

Thus, there are three elements in the SIMNET cost equation: the simulators, the gateways, and the long-haul network. In their work to date, DARPA and BBN have done a creditable job of achieving a balance among these three elements. The panel came to the conclusion that there was enough flexibility within this structure to support the anticipated growth in the operational requirements. There are many things that can be done to improve the long-haul efficiency. To cite a few: (1) the data elements interchanged by the simulators can be shortened (e.g., by not sending altitude information for land vehicles); (2) the filtering at the gateways can be made more sophisticated to further restrict the broadcast dissemination; (3) the data elements can be sent less often if more sophisticated dead-reckoning algorithms are used by the simulators. There are currently plans to add low-cost front-end communications processors to the individual simulators, thereby making more computing power available for the simulator processing at either transmit or receive end. Also, it would be relatively economical to increase the processing capabilities of the gateways. General improvements in computing capabilities that can be expected coupled with reasonable computing enhancements make us confident that the DARPA goals can be achieved.

E. INTEROPERABILITY AND STANDARDS

DARPA and BBN took the understandable point of view that meeting the operational requirements was of paramount importance. In so doing, they used nonstandard protocols because the available standard ones were not suitable. While some OSI purists may take a dim view of this, we on the panel felt that DARPA and BBN did the right thing in adopting the experimental DoD Internet ST and in using a special protocol on top of it that combines elements of several of the higher layers. ST should become an OSI standard and the current structure can evolve to a more standard variety. It was clear to us that BBN has paid attention to standards and has an evolutionary path to standardization.

We should never forget that protocol standardization is a means to an end, not an end in itself. Interoperability is important for simulators that have a need to interoperate, not for all of the simulators of the world that cannot interoperate and, what's more, have no need to do so.

F. CONCLUSIONS AND RECOMMENDATIONS

DARPA and BBN have done an exemplary job on a difficult and important problem. Their approach is not quite standard, but it works, it is scalable and can evolve to accepted standards. Despite our confidence in the basic approach, I think that it is important that DARPA continue to support R&D in the long-haul communications aspects of SIMNET, addressing both the short-term issues associated with the Army's near-term program and longer term issues dealing with a much-expanded network. While it has been prudent to avoid satellite connectivity thus far, there are operational situations in which satellites offer the only reasonable way of obtaining wide bandwidths (e.g., communications with ships at sea). I therefore feel that this continued R&D should include investigating the use of satellites for long-haul connectivity.

SUMMARY OF SIMNET PANEL REVIEW

David L. Mills
University of Delaware

A. INTRODUCTION

On 1-2 March 1990 a SIMNET program review was held at the Institute of Defense Analysis in Arlington, Virginia. The review panel, of which I was a member, met to hear technical briefings on the SIMNET program, specifically on the questions of scaling and protocol architecture for wide-area deployment. The following are my impressions on these and related issues.

My overall assessment of the technical direction of the SIMNET program is quite positive, with some reservations to be discussed later in this report. The BBN technical presentations were sound, the presenters well qualified, and the background documents most thorough. It is evident that the protocol architecture has been profoundly driven by the need to deliver timely performance in exercises involving up to thousands of simulation entities scattered over major portions of the globe. In some aspects these needs have overridden conventional wisdom that suggests conformance to the architectures and protocols being developed by the standards community. It is these issues that constitute the major thrust of this report.

B. APPROACH AND SCOPE OF THIS REPORT

SIMNET is fundamentally a large, distributed simulator system involving many fighting machines or entities that move in real time over a common three-dimensional terrain database. Some of these machines, perhaps only a small fraction, are controlled by real people--drivers, gunners, and pilots--others are semi-autonomous and controlled as a group by a commander, and still others are completely autonomous and function according to preprogrammed plan. Individual simulation exercises can involve the entire resources of the network or be split into autonomous groups of separate simulations. While the planners envisage a network dedicated to the simulation mission, it is expected that a certain

The panel was asked to assess the current state and future plans with respect to the following objectives:

1. Is the present technical approach appropriate for the anticipated growth over the next decade? What technical refinements will be necessary in the short term in order to provide for this growth?
2. What hardware/software/network capabilities will be necessary in order to achieve acceptable cost/performance? What showstoppers may exist in the short or long term that may seriously constrain the technical evolution?
3. BBN has targeted a specific model based on current technology. Is this likely to be changed in significant ways as new technology develops or existing technology and standards evolve in expected ways?

The remainder of this report discusses these issues in the context of network engineering, protocol engineering, and migration to ISO protocols. It does not discuss security issues. The report ends with a conclusions section based on panel discussions and subsequent analysis.

C. NETWORK ENGINEERING

The briefings included much discussion on the mission and architecture of the network. Obvious choices for implementation include the use of emerging civil networks versus DoD mission networks; whether the network itself would be general purpose or special purpose in nature; and whether the network is to have a defined service and dedicated application or whether the assets required could be shared with other applications.

The SIMNET architecture is presently based on a DoD-mission, special- purpose, dedicated-application architecture. Clearly, these run counter to conventional wisdom that says large, integrated, general-purpose networks provide the best performance at lowest cost. However, at the present stage of development, the SIMNET mission requirements for near-simultaneous multicast delivery preclude the use of conventional (e.g., X.25) technology and protocols. However, should these requirements become more universal--and such a case could be made for use in multiway real-time conferencing systems, both DoD and civil--there may be real merit in pursuing a more integrative approach. In fact, the mission requirements for a limited capability for real-time conferencing in exercise planning and review suggests that SIMNET technology may itself constitute a technology adaptable for general use.

In the SIMNET design the fundamental service required is the near-simultaneous delivery of position reports and strike reports of every entity simulated to every other entity throughout the network. This ordinarily requires a packet rate from a packet in several seconds to upwards of 10 packets per second. The service delay expected of the network must be less than a few hundred milliseconds for realistic exercise management. However, the requirements for reliability can be relaxed somewhat, since prediction techniques allow for some packet loss. BBN demonstrated several clever techniques which allow the more time-critical events such as missile firing and impact to be processed in the entity that causes them and then distributed throughout the network. In this case the reliability must be assured.

There is some question about the applicability of satellite technology to the SIMNET mission. The principal objection to the use of this technology is the 270-msec propagation delay inherent in the geosynchronous space segment. Considering the obvious applicability and unique cost effectiveness of satellite technology, I believe every effort should be made to thoroughly research and evaluate this issue. In principle, the use of VSATs and emerging technology represented by the NASA ACTS program could vastly simplify the protocols, reduce expense, and improve reliability on a global scale.

There was discussion concerning dynamically reconfigurable network resources, such as dial-up services. At present, such services are relatively expensive and probably do not represent a cost-effective alternative to dedicated circuits for most portions of the SIMNET service area. However, I believe there may be real potential in using dedicated facilities with digital automatic cross-connect (DAX) capabilities. Such systems would allow customer reconfiguration within the dedicated facilities leased from a common carrier and allow the network to be reconfigured for specific exercises.

D. PROTOCOL ENGINEERING

The key ingredient in the SIMNET architecture is the design of the simulation protocol. In many respects this is the principal distinguishing characteristic of SIMNET, as the other protocol functions required for the simulation mission can in principal be provided by off-the-shelf network technology. The principal requirements for the simulation protocol are as follows:

1. The protocol should provide a standard, transparent interface to lower level protocols, both for the IP and ISO suites. It must make as few assumptions

on the lower level services as possible in order to provide for as much application portability as possible.

2. The protocol must provide reliability tuned to performance requirements. In conventional models the metric used to assess reliability is static and almost never quantified. In the SIMNET application the reliability metric must be selectable within limits, depending on the particular service required, best-effort (reserved) multicast by dead-reckoning and repetition, assured multicast for strike reports and conventional assured monocast for event logging and network management.

3. The protocol must incorporate association management to define multiple distinct subnets and to allow entities to join and leave the simulation as required by the exercise plan.

4. The protocol data unit (PDU) encoding used must be fast and efficient, as well as adaptive to handle editing on-the-fly by intelligent routers. The PDUs must be readily aggregable for efficient encapsulation and transmission in order to reduce header overhead at the lower levels.

It was apparent at the briefings that the presenters were acutely aware of the impact of performance expectations on the engineering design. BBN has participated for many years in the DARPA-funded Internet program, so the design approaches found useful in that program could be expected to appear prominently in the SIMNET design. There are two technologies that appear as drivers: a protocol architecture based on datagram principles, and a network architecture supporting a semi-reliable, reserved-resource multicast service.

In low-latency LANs there is usually little concern for transmission delay since the bandwidths on the media are usually large compared with the offered traffic, and delays are usually small. On LANs where bandwidth is not a premium, a multicast function is easily achievable using any of several technologies, such as Ethernet. The present SIMNET design simply encapsulates the SIMNET PDU as an Ethernet packet and broadcasts it on the local Ethernet.

However, in the case of WANs, bandwidths are usually much less than Ethernet (10 Mbps) and delays are usually dominated by queueing delays in packet switches. The BBN designers have concluded that on LANs a resource-reservation protocol is required in order to provide strict control over delays. The BBN approach for WANs borrows conveniently from the DARPA WIDEBAND system originally developed for satellite use, but now adapted for terrestrial use in DARPA testbed networks. This technology has stood the test of time, having been in use for several years as a vehicle for distributed multimedia conferencing for DARPA meetings, for example.

The reservation/multicast protocol chosen is called ST and was designed some years ago for multicast real-time speech applications. For ease of integration the designers have chosen to encapsulate SIMNET PDUs in ST packets and ST packets in IP datagrams. The question is, can a case be made for the use of a resource-reservation multicast protocol in general and ST in particular as a SIMNET requirement? There are alternatives to the use of ST, including IP multicasting, which is rapidly becoming a ubiquitous feature of the Internet. However, while IP multicasting includes an association-management function, it does not include a resource-reservation function, which I believe should be more an attribute of the network itself than the protocol used in the entities and routers.

It is my conclusion that the use of IP multicasting together with special-purpose resource-reservation services built into the network itself have not been adequately considered. It may be that the encoding economy provided by the compact ST header may not be much effective, unless something like IP header compression were employed as well. Certainly, resource reservation techniques within the network itself have been proposed and studied previously, most notably at BBN. The conventional wisdom coming from the Internet engineering study groups is that the use of these techniques should become much more widespread as network technology continues to evolve and flourish. Therefore, an appropriate strategy might well be to expect the resource-reservation function to be exercised in the network (routers) and the multicast-setup function to be exercised as part of the network-layer functions in the hosts and routers.

An approach encapsulating SIMNET PDUs directly in IP datagrams with IP multicast group addresses and calling on specific network services for routing and delay control has the advantages of near ubiquity, network independence (assuming intrusive routers, which may be required anyway) and simple migration to ISO Connectionless Network Protocol (CLNP), should that become viable. It may be that ST represents a viable vehicle to provide resource reservation on specific LANs, but there could be other mechanisms as well. It may even be possible through clever engineering to adapt the IP multicast address for use directly as a stream identifier for ST. The goal is to make ST a feature of the network, not a feature of the protocol.

E. ADAPTIVE COMPRESSION AND ROUTING

The real-time nature of the SIMNET application invites some interesting optimization techniques which provide near simultaneity while relaxing the need for the shortest transmission delays. An example described by BBN occurs when a firing command is issued by a particular entity. The command is executed and the impact predicted at the instant of firing, followed immediately by a strike report broadcast to all entities. This gives some time for all entities to receive the broadcast and compute the effects on the graphics display at the instant of the predicted strike. As most simulation messages must be delivered to many other entities, there is the potential to overwhelm network resources, especially when these entities are scattered over a global network.

Adaptive compression recognizes that not all entities need receive the full precision and data rate to maintain a realistic simulation. This allows an important degree of compression and more efficient use of network links. There are three ways this can be done: adaptive encoding, adaptive refresh rate, and adaptive routing. All three require intrusion of the router function at the upper layers of the protocol stack.

One technique used to reduce the level of traffic is to estimate vehicle position using past coordinates and velocity. This provides a robust position estimate while reducing the packet rate and impact of lost packets. From the data presented at the briefing, BBN has evaluated the technique and engineered refresh rates and dynamic adjustments appropriate for each vehicle. This might be termed transmitter-directed refresh rate. However, there may be additional benefit to be gained by considering not only the required refresh rate for precision location (in the order of a meter for the M1 tank), but also the effective resolution on the part of a distant observer. This might be termed receiver- directed refresh rate. While not mentioned in the briefing, these ideas can be extended to affect the rate at which individual entities receive updates; some may not need to be updated as often as others. Another way to reduce the traffic rate might be termed intelligent obfuscation or adaptive encoding. Distant observers may not need the full precision capability of the full PDU format; intelligent routers may compress the data by truncating the low-order bits of the position information and repositioning the transformation coordinates. One obvious thing to do would be to use nonlinear transformation coordinates, in which the resolution varies along the axes proportional to the expected trajectory error. It is not clear to what extent BBN may already be doing this, but there may be considerable merit in pursuing these issues with considerable vigor.

An intelligent routing algorithm might intentionally discard packets for certain entities by computing the intersection between the direction vector and possible obstructions between one vehicle and another. This might be considered the limit of the adaptive encoding process when the number of bits required drops to zero. However, adaptive routing also recognizes that not all entities can "see" all other entities and therefore the spanning tree can be edited accordingly.

All three of the above techniques require knowledge in the routing and forwarding functions of intricate details of the terrain database and entity geometry. Performing them effectively may require substantial CPU resources, which may result in performance penalties in throughput. I conclude that an exceptional degree of ingenuity is required and likely a substantial investment in prototyping, testing, and refinement.

F. MIGRATION TO ISO PROTOCOLS

At the briefing there was much discussion of the impact of the standards process on the future evolution of SIMNET. Specifically, the issue of conformance to ISO protocols was raised in the context of procurement and interoperability. The primary drivers for this include the opportunity to facilitate peer review, international coordination (especially NATO), and multivendor implementation agreements. It was pointed out that, while these issues are important drivers, it is more important to standardize the network services than the application itself. And, in fact, it is more important to standardize at the periphery of the system than internally. This follows the recent trend in application-level routers.

I believe that even at the network-services layers the urge to rush to standardize should be resisted. The ISO study groups have only begun to address the issue of multicasting; this issue and other engineering issues concerned with efficiency and reliability have not been dealt with effectively by ISO study groups in the past and are not likely to be so in the future. I conclude that the BBN approach, which emphasizes pragmatic engineering with later development and migration to the standards process, is most appropriate.

There are some areas for which a case can be made for standardization even now, specifically PDU encoding and association management. In particular, the PDU fields and encoding could be adapted to conventional ISO principles, specifically ASN.1 encoding. There may be danger in this approach. The encoding/decoding overhead can be extreme and, as BBN reports, the CPUs are already under strain at the present speed and size of the system. In fact, one panelist reported rates of 50 PDUs per second on a Sun/3 with

ASN.1. While this dismal performance may be attributed to experiment inefficiencies and while special-purpose hardware is likely to become available, it nevertheless underscores the importance of performance as a critical issue to SIMNET, but not necessarily to the standards community.

It was suggested that the association management function provided by the Association Control Service Element (ACSE) facility could be adapted for use by SIMNET. However, the ISO principles are intended primarily for point-to-point applications and unsuited in present form for multicast with a large number of destinations. These observations are not necessarily showstoppers, but they do suggest that considerable investment in engineering analysis, design, and promotion within the standards community will be necessary before a robust set of standards can be developed for SIMNET.

One of the chief objections to the particular BBN design is that it includes no clearly defined presentation, session, or transport layers with respect to the ISO model. BBN's answer to these objections is that interfaces could be developed if and when the appropriate application and network layer services became available. I regard the objections as wholly a red herring and agree that BBN should concentrate on the engineering development of a robust, performance-oriented product and leave the standards and interface details for future study.

The issue of network management is much more clear to me. BBN is using a proprietary protocol developed some time ago for use in several of their network products. There appears to me no reason other than expediency why recently standardized protocols and service definitions such as SNMP could not be used instead. In fact, this might be a relatively simple thing to do, and it might cost very little.

G. CONCLUSIONS

An index to the future scalability of SIMNET can be estimated from the number of packets per second that can be handled by a single simulation entity, in this case a single CPU and LAN interface. The present system tops out at about 1,000 entities, which produce on the order of 800 packets per second (pps). The panel concluded from the briefings that the present SIMNET architecture and protocols can be scaled upwards by an order of magnitude in entities probably without changing the architecture or protocols, due to anticipated developments in CPU speed and transmission costs. Such a system would have on the order of 2,000 entities, which would generate on the order of 3,200 pps. It was suggested that a front-end communications board or faster CPU (e.g., 68040) could handle this using the same software. A faster network would be necessary, as well as careful engineering to avoid congestion and excessive delay. I concur with this view and emphasize that satellite technology is ideally suited for the network technology if problems in latency can be overcome. According to the objectives stated in the introduction of this document, this course would be appropriate through the year 1992.

The panel concluded that scaling another order of magnitude beyond that is likely to stretch the current architecture and protocols and may require re-engineering to achieve it. Four times the present number of entities requires 16 times the basic rate or 12,800 pps. Present-day network routers can handle such rates, but only using specially engineered interfaces and memory ports. Also, at these rates the capacity of Ethernets begins to wilt, so even the familiar LAN technologies run low on steam. I concur with this view and believe such a course is possible with full understanding that it is a dead end and not likely to evolve beyond this order of magnitude. This course may satisfy the objectives through 1994.

However, in order to scale the entities up by a factor of 10 or more, which is the strawman objective for the end of the century, the panel concludes that substantial changes in the architecture and protocols will be necessary. In particular, improved position-report, adaptive-compression, route-filter, and resource-reservation algorithms will be necessary and will require further specialization and distance from the standards process. I concur with this view and emphasize that the network technology will likely be based on high-speed protocols and fiber technology. Interfaces will necessarily be highly specialized to the application and contain considerable intelligence to offload the CPU, which may itself involve RISC architecture and include special-purpose graphics engines. In short, the development strategy may be optimized for further specialization and away from a combined mission and standardization process.

APPENDIX B
BIOGRAPHICAL SKETCHES OF PANEL MEMBERS

APPENDIX B
BIOGRAPHICAL SKETCHES OF PANEL MEMBERS

JEFFREY D. CASE

Jeffrey D. Case is an Associate Professor of Computer Science, Department of Computer Science, at the University of Tennessee. He teaches, conducts research in networking and network management, and directs the Computer Science Laboratories. He received the Ph.D. degree from the University of Illinois in 1983. He formerly served as the Chief Engineer of the University of Tennessee Computing Center where he was responsible for local area networking, wide area networking, engineering, VAX/VMS systems, and Unix systems and workstations including responsibility for a large state-wide network. Before moving to Tennessee, he served as a faculty member and administrator (computer services) at Purdue University.

His current research is in the area of networking, network protocol design, and network management. He is a coauthor of the Simple Network Management Protocol, the network management protocol for TCP/IP-based internets and has published several papers and spoken widely on the subject. His expertise and activities in network management led to his being named a Martin Marietta Professor of Computer Science and recognition as the "Newsmaker" of the month by Data Communications magazine in January, 1990.

Dr. Case is the author of one of the leading vendor-independent reference implementations of the SNMP. This implementation has found wide acceptance by network users as well as communications and computer hardware, software, and service providers, both domestic and international.

DANNY COHEN

Danny Cohen has a B.Sc. in Mathematics from the Technion in Israel and a Ph.D. from Harvard. He is currently the Director of the Systems Division at the Information Sciences Institute.

In 1967 he developed the first ever real-time visual flight simulation on a general purpose minicomputer. In 1968 he developed the first ever all-digital mass land radar simulation. Since then he has been involved in real-time computer communication. In the early 1970s he introduced the concept of real-time protocols both for the ARPAnet and for the Internet. He is currently active in internet research.

Dr. Cohen is a bonafide member of the Flat Earth Society.

DALE B. HENDERSON

Dale Henderson joined the Los Alamos National Laboratory in 1966 upon completion of his Ph.D. at Cornell University. After four years in experimental plasma physics, he moved to the (then new) laser induced fusion program. In 1975 he became leader of the theory group in that project. In 1979 he moved to project management of computer code development for the nuclear weapons design program. Soon after President Reagan's "Star Wars" speech, he recognized the need for a flexible comprehensive simulation model and began the DETEC (Defensive Technology Evaluation Code) project at Los Alamos. DETEC was adopted as the major software vehicle at the SDIO'S National Test Bed (NTB) in 1988. Having served the NTB Joint Program Office from its beginning, Dr. Henderson undertook an FY 1990 assignment to the SDIO as Chief Scientist of the NTB.

IRWIN L. LEBOW

Irwin L. Lebow has been an independent consultant since 1987. He spent most of his career at MIT's Lincoln Laboratory, leaving in 1975 to join the Defense Communications Agency (DCA) as Chief Scientist-Associate Director, Technology. He left DCA in 1981, serving successively as Vice President, Engineering, at American Satellite Co. (now Contel ASC) and Vice President at Systems Research and Applications (SRA) Corp. He earned the B.S. and Ph.D. degrees in physics at MIT.

At Lincoln Laboratory he contributed to the design of one of the first all-solid-state computers and later to Lincoln's pioneering work in satellite communications. At DCA he had responsibility for the Agency's RDT&E program and was centrally concerned with the upgrade programs both in the Defense Communications System and in the World Wide Military Command and Control System. At American Satellite he led both the day-to-day engineering and R&D efforts of the company that pioneered digital satellite communications in the commercial marketplace. He was responsible for both commercial and military communications efforts at SRA, where he has maintained his association as a Senior Consultant.

He is a member of the DCA Scientific Advisory Group and the Radio Engineering Advisory Committee of the Voice of America. His awards include the Defense Department's Meritorious Civilian Service Medal. He is a fellow of the American Physical Society and the IEEE. He has authored many papers, co-authored Theory and Design of Digital Machines (McGraw-Hill, 1962), and authored the chapter on satellite communications in Digital Communications (T.C. Bartee, ed., Sams, 1986). His latest book, The Digital Connection, is to be published by Computer Science Press/W.H. Freeman & Co. in 1990.

DAVID L. MILLS

David L. Mills is Professor of Electrical Engineering at the University of Delaware and presently leads projects in high-speed networks and internetworking research sponsored by the Defense Advanced Research Projects Agency (DARPA) and National Science Foundation (NSF). His research activities have been concentrated in the areas of network architecture, protocol engineering, and experimental studies using the DARPA/NSF Internet system. He is a member of the Internet Research Steering Group and formerly chaired the Internet Architecture Task Force. He is also an advisor to the NSF and was principal architect of the NSFNET Phase-I Backbone network.

Before joining the Delaware faculty in 1986, Dr. Mills was a Director (Networks) at M/A-COM Government Systems Division (Linkabit) and led DARPA-sponsored R&D projects in packet-switching network architectures and application protocols. Before that he was a Senior Research Scientist at COMSAT Laboratories, where he worked in the areas of packet-switching satellite and internetworking technologies, and Assistant Professor of Computer Science at the University of Maryland, where he worked on several research projects in distributed computer networks and operating systems.

Dr. Mills earned a Doctorate in Computer and Communication Sciences at the University of Michigan in 1971 and has held postdoctoral positions at the University of Edinburgh (Scotland) and U.S. Defense Communications Agency. He has published and lectured extensively on data communications, computer networks, and operating systems and has been a consultant to a number of corporations and government agencies. He is a member of Sigma Xi, Association for Computing Machinery, and IEEE Computer Society.

APPENDIX C
READ-AHEAD DOCUMENTS

APPENDIX C
READ-AHEAD DOCUMENTS

Garvey, R.E. and Radgowski, T. Data Collection and Analysis: The Keys for Interactive Training and Combat Readiness. Cambridge, MA: BBN Systems and Technology Corporation, (Undated).

Kanarick, C.M. and Pope, A.R. Summary of SIMNET Protocol Changes for the Period August 1989 through January 1990. Cambridge, MA: BBN Systems and Technology Corporation, 12 January 1990.

Pope, A.R. The SIMNET Network and Protocols (Report 7102). Cambridge, MA: BBN Systems and Technology Corporation, July 1989.

Pope, A.R. and Miller, D.N. The SIMNET Communications Protocol for Distributed Simulation. Proceedings of the Sixth Annual Technology in Training and Education (TITE) Conference. Biloxi, MS: March 1988.

SIMNET Local Area Network (LAN) Overview. Cambridge, MA: BBN Systems and Technology Corporation, 9 January 1990.

SIMNET Long Haul Network (LHN) Overview. Cambridge, MA: BBN Systems and Technology Corporation, 9 January 1990.

SIMNET Overview. Cambridge, MA: BBN Systems and Technology Corporation, 9 January 1990.

Terrestrial Wideband Network (TWBnet) Overview. Cambridge, MA: BBN Systems and Technology Corporation, 9 January 1990.

Wever, P., Lang, E., and Smyth, C.S. SIMNET Database Interchange Specification (Report 7108). Cambridge, MA: BBN Systems and Technology Corporation, July 1989.

The following papers were distributed at the meeting in support of the Martin Marietta/SSDS presentation:

Sabo, L.M. Distributed Simulators Architecture (DSA). Littleton, CO: SSDS, Inc. 26 January 1990.

Sabo, L.M. Interactive Simulation Protocol (ISP). Littleton, CO: SSDS, Inc. 31 January 1990.

Sabo, L.M. Transport Layer Protocol Options for the Distributed Simulators Architecture (DSA). Littleton, CO: SSDS, Inc. 13 February 1990.

Sabo, L.M. Using the Simple Network Management Protocol (SNMP) within the Distributed Simulators Architecture (DSA). Littleton, CO: SSDS, Inc. 23 February 1990.

APPENDIX D
REVIEW MEETING ATTENDEES

APPENDIX D
REVIEW MEETING ATTENDEES

PANEL MEMBERS:

Professor Jeffrey D. Case
Department of Computer Science
107 Ayres Hall
University of Tennessee
Knoxville, Tennessee 37996
615-974-0822
615-573-1434
Case@CS.UTK.EDU

Dr. Danny Cohen
Systems Division
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, California 90291
213-822-1511
Cohen@ISI.EDU

Dr. Dale B. Henderson
Chief Scientist
SDIO National Test Bed
SDIO/PTN (NTBJPO)
Falcon AFB, Colorado 80912
719-380-3230
DBH@LANL.GOV

Dr. Irwin L. Lebow
2800 Bellevue Terrace NW
Washington, DC 20007
202-333-1836

Professor David L. Mills
Department of Electrical Engineering
University of Delaware
Newark, Delaware 19716
302-451-8247
Mills@UDEL.EDU

BOLT BERANEK AND NEWMAN:

Steve Blumenthal
Jerry Burchfiel
Alexander McKenzie
Duncan Miller
Arthur Pope
Claudio Topolcic
Rolland Waters

BBN Systems and Technology
33 Moulton Street
Cambridge, Massachusetts 02138
617-873-2000

MARTIN-MARIETTA/SSDS:

Michael Sabo
SSDS, Inc.
Suite 200
1101 W. Mineral Avenue
Littleton, Colorado 80120
303-798-5520

DEFENSE ADVANCED RESEARCH PROJECTS AGENCY:

LTC (P) James Shiflett
LTC Mark Pullen
Dr. Ira Richer

Information Science and Technology Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209-2308
202-694-5800

ADDITIONAL ATTENDEES:

John Collins
US Army PMTRADE
ATTN: AMCPM-TND-EC
12350 Research Parkway
Orlando, Florida 32826-3276
407-380-8077

Dexter Fletcher
Institute for Defense Analyses
1801 North Beauregard Street
Alexandria, Virginia 22311
703-578-2837
Fletcher@A.ISI.EDU

VADM Leland S. Kollmorgen (USN, Ret)
1902 Joliette Court
Alexandria, Virginia 22307
703-528-8901

David Wood
MITRE Corporation
7525 Colshire Drive
McLean, Virginia 22102
703-883-6000

APPENDIX E
LONG-HAUL NETWORK BRIEFING PRESENTED BY BBN

SIMNET Overview

Duncan C. Miller, Sc. D.
BBN Advanced Simulation Division
33 Moulton St., Cambridge, MA 02138
(617)-873-3334
Internet: Miller@bbn.com

15 January 1990

"Distributed simulation" approach

- An object-oriented simulation architecture
- No central computer is used for event scheduling, etc.
- Each simulation microcomputer is autonomous
  - *responsible for maintaining state of one simulation element*
  - responsible for sending messages to others, as necessary
  - responsible for interpreting and responding to messages
- As network expands, each new simulator brings its own resources
- Each simulator has its own copy of the non-changing world
- Simulators communicate only <u>changes</u> in world state
- "Dead reckoning" used to reduce communications processing

Dead reckoning technique

- Each simulator maintains a simplified model of the state of all others
- Last reported states are extrapolated until new update arrives
- Each simulator must maintain a dead reckoning model of itself
- Must broadcast update whenever true state diverges from extrapolation
- Update interval depends on actions of vehicle crews
- Minimum update interval is *currently* 1/15 second; max is 5 seconds
- Updates contain externally "visible" information
  - position, velocity, attitude
  - turret azimuth (if any), gun elevation, etc.
  - dust clouds, smoke column, muzzle flash, etc.
  - *current maximum is 256 bytes (2048 bits) per update*
- State update protocol is self-healing
  - if an update is missed, extrapolation of old state continues
  - next update corrects state and initiates new extrapolation

Balancing communication, computation, and precision

- Dead reckoning approach involves design tradeoffs
  - communications load is reduced (typically by factor of 10)
  - computational load is increased (and grows with size of net)
  - communications load increases as thresholds are reduced
- Higher-order dead reckoning algorithms change balance point
  - higher-order extrapolation can stay within threshold longer
  - communications processing load will be reduced
  - dead reckoning computational loads will be increased
- Currently using first-order extrapolation for all vehicles
  - average rate of 1 update/sec per ground vehicle
  - average rate of 3 updates/sec per air vehicle
- With second-order extrapolation, air vehicle drops to 1 update/sec
- Ground vehicles don't drop much (because of ground interactions?)

Local Area Network (LAN)

- SIMNET protocols do not depend on what network is used
- Ethernet selected because of low cost, multiple vendors
- Ethernet *interfaces are available for virtually every computer*
- Up to 1024 devices can be connected to a local area net
- A 1000-vehicle exercise is well within efficient operating range
- At this point, however, remote vehicle processing begins to dominate simulation hosts
- Other LANS are being developed, including FDDI

Long-haul links

- Initial network used Butterfly gateways, 56 kilobit/second dial-up lines
- 56 kilobit/second lines support 100 vehicles (each direction)
- Multiple lines can be used to increase capacity
- Dedicated T1 lines (1.544 Mbps) can also be used
- Long-haul links now carrying digitized voice traffic as well as data
- Now using Terrestrial Wideband Net with ST Protocol
- TWBnet also provides DoD Internet Protocol service for file transfers, etc.

Terrestrial Wideband Network

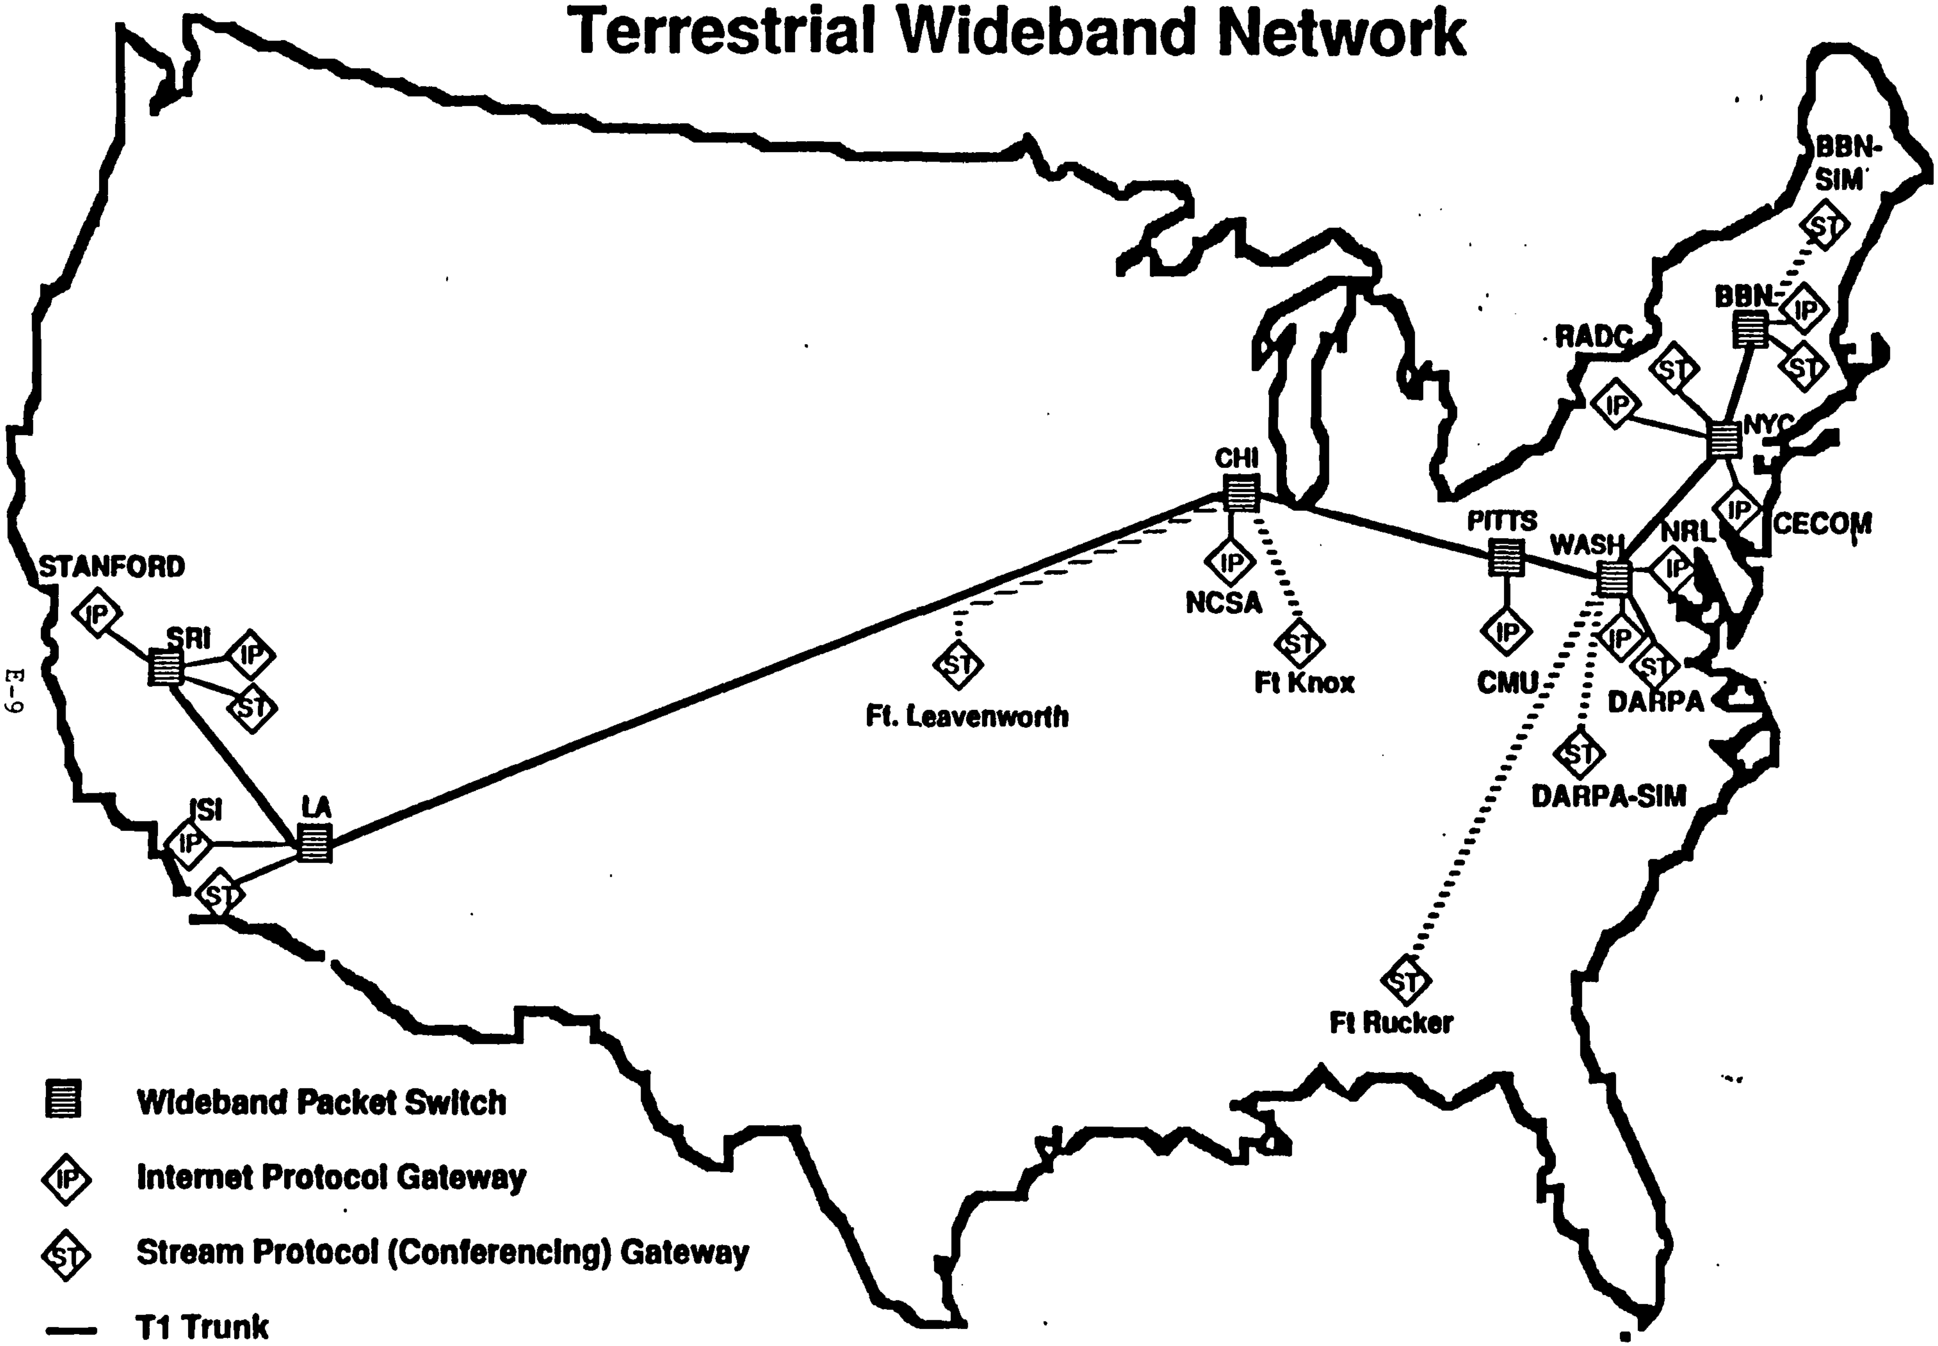

Types of simulations supported by SIMNET

- Manned vehicle simulations
- Automated (support) simulations
- Semi-automated simulations

Manned simulators

- Provide full-crew control/display interactions
  - M1 Abrams Main Battle Tanks
  - M2/M3 Bradley Fighting Vehicles
  - Generic fixed- and rotary-wing aircraft
  - Generic air defense artillery

Automated simulations

- Are not intended to represent crew control/display interactions
- Represent higher-level resource allocation functions that are critical to combined arms operations
  - Fire support (howitzers, mortars)
    - fire support officer enters firing missions
    - if battery is within range and available, mission is executed
    - missions are also limited by resupply rates, etc.
  - Logistics (fuel, ammunition trucks)
    - S-4 controls and dispatches vehicles
    - vehicles "teleport" across battlefield with realistic delays
    - vehicles must return to supply points for reloading
  - Maintenance and recovery
    - battalion maintenance officer controls and dispatches vehicles
    - vehicles "teleport" across battlefield with realistic delays
    - repair actions consume realistic time intervals
    - if correct actions are performed, simulated damage is repaired

Semi-automated simulations

- Unlike automated simulations, involve realistic motion over terrain
- Commander provides goals and objectives, may designate routes
- Vehicles move in formation, avoid obstacles, follow terrain
- Vehicles scan for enemy, make target engagement decisions
- Commander may override priorities, *redirect units at any point*
- Semi-automated vehicles are indistinguishable from manned vehicles
- Current forces implemented (each with red and blue versions)
  - armor
  - mechanized infantry
  - helicopter
  - fixed-wing aircraft
  - air defense artillery

Data collection and analysis system

- Data Loggers record all state update messages for later replay
- Replay provides "*time travel*" capabilities for observing action
- Plan View Display provides continuous view of battlefield state
  - controls include pan, zoom, intervisibility plots, etc.
  - replay controls include fast time, rewind, freeze frame, etc.
- DataProbe and RS/1 provide statistical analysis and plotting tools

Stealth vehicle

- Invisible to other participants (can be used during exercise)
- Used primarily for after-action reviews
- Very simple, "flying carpet" dynamics
- *Can teleport to any desired location in the simulated world*
- "Tractor beam" can be used to attach to, and follow, any vehicle

Electro-optical and thermal imaging systems

- Currently using simplified representation of sensor characteristics
- Same line-of-sight algorithms used as for daylight displays
- Modified texture maps used to simulate appearance of terrain
- Using alternative models to simulate appearance of targets
- Bank-switching added to CIG to permit switching of models

Digital data communications among vehicles

- Simulating proposed Inter-Vehicle Information System displays
- Own-vehicle position displayed, plus other vehicles in unit
- *Situation reports, etc., can be composed and transmitted*
- Target locations can be transmitted among displays
- Target assignments and other orders can be received

Radar simulation

- Local horizon continuously computed, target line-of-sight determined
- Detection probability depends on target aspect, range, masking
- Radar units themselves broadcast "appearance packets", including which sectors are being illuminated
- Radar warning receivers calculate energy received, set off alarm
- Radar emitters are vulnerable to radiation-seeking missiles

Missile models

- Missiles are modelled in three broad categories
- Optically-tracked missiles
  - hit depends primarily on tracking accuracy of controller
  - missile flyout model runs in controller's simulator
  - visual effect messages broadcast by controller's simulator
- Fire-and-forget missiles
  - hit depends primarily on evasive actions of target
  - missile flyout model runs in target vehicle's simulator
  - missile appearance messages broadcast so everyone sees trajectory
- Laser-designated missiles
  - success requires line-of-sight between designator and target
  - also requires line-of-sight between missile and target
  - target simulator informed when it is being designated
  - target vehicle adds designation codes to its appearance messages
  - missile flyout model calculates line-of-sight, broadcasts trajectory

Radio communication simulation

- Voice signals digitized at transmitter, transmitted via network
- Signals digitally compressed to reduce bandwidth requirements
- Header information added to describe
  - Transmitter location, and frequency
  - Antenna orientation and emitted power
- Receiver simulators compute which signals are heard
  - VHF *signal strength computed from* path length to source
  - path length includes diffraction across major obstacles
  - receiver model determines which signals are captured (FM radios lock in on strongest signal, for example)
- Only "winning" signal gets decoded, mixed with noise
  - this signal may be a jammer
  - jamming may be intentional or unintentional
- Direction finders can also use emitter location information
  - can be used by own forces, for navigation
  - can be used by enemy anti-radiation weapons

Mine laying and clearing

- Automated simulation, running on Macintosh workstation
- Minefields laid in designated areas (which takes realistic time)
- Server computes vehicle proximity to mines, announces detonations
- Each simulator computes possible damage, as for artillery round
- Minefields can be cleared in designated areas (which also takes time)

SIMNET program status

- SIMNET program is currently being concluded
- 236 simulators are operational at eleven sites
- Nine sites being transitioned to Army under competitive COMS contract
- Two sites will remain under DARPA control for further development
- Competitive contract being awarded for Advanced Distributed Simulation
- ADS contract will explore
  - very large scale wargaming (10,000 platforms or more)
  - intelligent networking approaches to support such wargames

BIBLIOGRAPHY

Burchfiel, Jerry, and Stephen Smyth. *Use of Global Coordinates in the SIMNET Protocol White Paper ASD-90-10*. BBN Systems and Technologies Corporation Cambridge, Mass., January 1990.

Downes-Martin, Stephen. *Replacing the Exercise Controller with the Enemy: The SIMNET Semi-Automated Forces Approach.* BBN Report No. 7221. To be presented at the International Training Equipment Conference, UK, April 1990. BBN Systems and Technologies Corporation Cambridge, Mass., December 1989.

Friedman, Dan, and Varda Haimo. *SIMNET Ethernet Performance*. BBN Report No. 6711. BBN Communications Corporation Cambridge, Mass., January 1988.

Garvey, Richard E., and Thomas Radgowski. *Data Collection and Analysis: The Keys for Interactive Training for Combat Readiness*. BBN Systems and Technologies Corporation Fort Knox, KY., April 17, 1989.

Miller, Duncan C., Arthur R. Pope, and Rolland M. Waters. *Long-Haul Networking of Simulators*. Proceedings of the 10th I/ITSC Conference. BBN Systems and Technologies Corporation Cambridge, Mass., November, 1988.

Pope, Arthur R., and Duncan C. Miller. *The SIMNET Communications Protocol for Distributed Simulation*. Proceedings of the 6th TITE Conference. BBN Laboratories Incorporated Cambridge, Mass., March, 1988.

Pope, Arthur R. *The SIMNET Network and Protocols*. BBN Report No. 7102. BBN Systems and Technologies Cambridge, Mass., July 1989.

Radgowski, Thomas, and Richard E. Garvey. *SIMNET-D: Combat Modeling Through Interactive Simulation*. Presented at the 57th MORS Symposium, June 1989. BBN Systems and Technologies Corporation Cambridge, Mass.

Stanzione, Thomas. *Terrain Reasoning in the SIMNET Semi-Automated Forces System*. BBN Report No. 7140. Presented at the Geo '89 Symposium on Geographical Information Systems for Command and Control, SHAPE Technical Center. BBN Systems and Technologies Corporation Cambridge, Mass., October 1989.

Wever, Peter, Eric Lang, and C.S. Smyth. *SIMNET Database Interchange Specification*. BBN Report No. 7108. BBN Systems and Technologies Corporation Cambridge, Mass., July 1989.

To obtain copies of these documents, please contact:

Ms. Cheryl L. Walsh
BBN Advanced Simulation Division
33 Moulton Street
Cambridge, MA 02138
(617) 873-3211

Overview

- "The SIMNET Network and Protocols", July 31, 1989
- requirements: the simulated world
- architecture: networks, layered protocols
- communicating vehicle appearance information
- supporting networks
- survey of protocol interactions
- representation of information
- communication compatibility
- *future work*

Everything the simulated world includes

- a region of terrain
  - typically tens or hundreds of kilometers on a side
  - populated with features: hills, rivers, roads, trees, buildings...
  - static — not changing in the course of a simulation
- a particular date and time
- vehicles that move dynamically and engage in combat
- supplies of munitions, such as fuel and ammunition
- the transfer of munitions from one vehicle to another
- weapons fire and its effects upon vehicles
- damage to vehicles and vehicle breakdowns
- repairs performed by one vehicle on another
- radar emissions and detection by radar

Goals of the SIMNET protocols

- a real-time network of hundreds of simulators
- ensure a consistent view of the simulated world
- be parsimonious and *efficient*
- allows efficient distribution of computation tasks
- be robust (*not error-sensitive*; self-correcting, if possible)
- easily accomodate new kinds of vehicles, weapons, phenomena...
- make available information useful for analysis

Architecture of the distributed simulation

- an internet of LANs and WANs

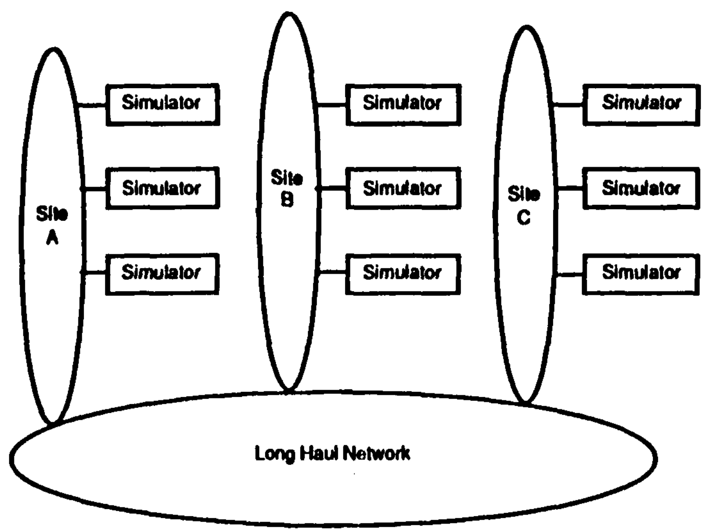

Varieties of simulators

- a simulator may take any of various roles:
  - simulate a single vehicle (e.g., a flight simulator)
  - simulate a group of vehicles (e.g., Semi-Automated Forces)
  - play a role in initializing other simulator (e.g., MCC system)
  - give a "God's–eye view" into the simulated world (e.g., Plan-View Display)
  - make an historical record (e.g., Data Logger)

Protocols

- three simulator-to-simulator protocols are defined:
  - a *simulation protocol* for representing the simulated world
  - a *data collection protocol* to support analysis
  - an *association protocol* to convey the other two

Layering of protocols

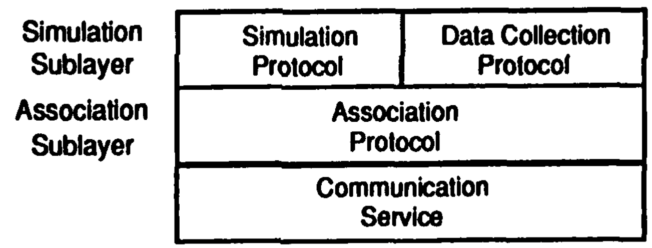

•

• the association protocol provide common services

Relationship to the OSI reference model

- protocols are defined within the framework of the OSI reference model

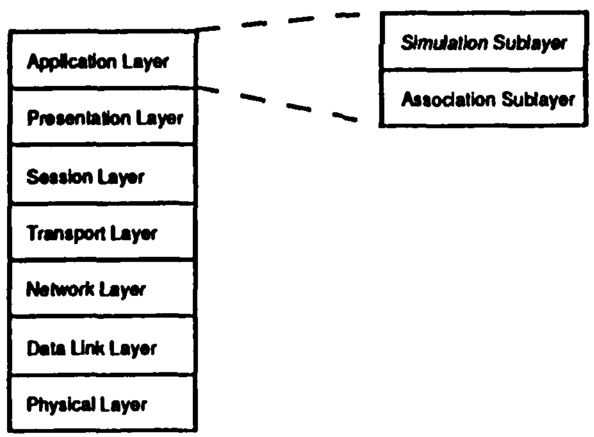

Distributed simulation concepts

- exercise: a simulated world maintained by some participating simulators
- exercise identifiers distinguish concurrent but independent exercises
- a simulated world is populated by vehicles
- each vehicle has:
  - static attributes: side, type, vehicle identifier
  - dynamic appearance: location, orientation, emissions
  - internal state: operational status, onboard supplies
- each vehicle's appearance is periodically reported with a state-update message

Communicating vehicle appearance

- each state update message is transmitted to all participating simulators
- most network traffic consists of state updates
- dead reckoning reduces the need for communication bandwidth
- various dead reckoning approaches are possible:
  - no use of dead reckoning
  - location updated using velocity
  - velocity updated using linear acceleration
  - rotation updated using rate of rotation
  - velocity updated using rotation

Effect of dead reckoning on network traffic

- dead–reckoning a tank using velocity only:

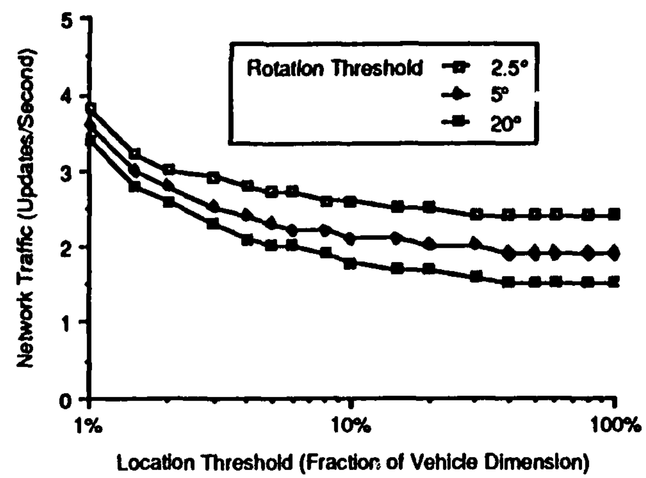

Effect of dead reckoning on network traffic

- dead–reckoning an aircraft using velocity, linear acceleration, and rate of rotation:

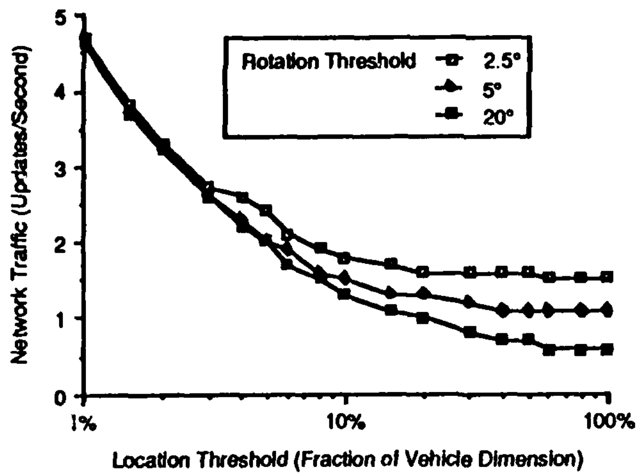

Current SIMNET dead reckoning

- vehicles are classified, partially according to dead reckoning method
    - *static class* — those that remain stationary
    - *simple class* — location updated using velocity
    - *tank class* — like simple, but has a turret
- discrepancy thresholds determine when state updates are issued
    - any discrete change in appearance (e.g., catching fire)
    - translation by 10% of vehicle's dimension
    - rotation about any axis by 3 degrees
    - movement of turret or gun barrel by 3 degrees

Data communication requirements

- SIMNET protocols are application layer protocols
- they are supported by an underlying network service:
- network must support broadcasting or multicasting of datagrams
- datagrams range up to 256 octets; most are 128 octets
- guaranteed delivery not required; occasional failures tolerated
- a level of performance determined by the "size" of the simulation
- various network technologies may be used
- network may be a combination of LANs and WANs
- Ethernet has been used succesfully as LAN

Network performance

- most network traffic is due to vehicle state updates
- traffic depends on number, type, and behavior of vehicles
- ground vehicles (tanks) produce an average of one update per second
- close-support air vehicle produce an average of six per second
- each update is communicated as a *128-byte datagram*
- each update must be communicated to all simulators in "real time"
- network delay, and delay variance, can be detrimental
- how much delay is acceptable depends on application:
  - relatively slow-moving ground vehicles can tolerate 300 ms
  - high–speed aircraft flying in formation cannot

Association protocol

- streamlined composite of certain transport, session, and application layer services
- eliminates need for separate transport and session layer protocols
- supports two modes of communication:
  - datagram service provides best–effort delivery
  - transaction service pairs request and response, provides retransmission
- clients are addressed by site number, simulation number
- clients belong to multicast groups

Simulation protocol

- includes procedures for:
  - activation of vehicles
  - deactivation of vehicles
  - vehicle state update
  - weapons fire
  - collision between vehicles
  - transfer of munitions between vehicles
  - repairs by (the crew of) one vehicle to another

Data collection protocol

- includes procedures for:
  - reporting the status of a vehicle or simulator
  - reporting various kinds of noteworthy events

Data representation

- aim is to minimize protocol dependence on machine architecture and language
- conservative restrictions enforced on data element alignment and size
- formal notation used: *data representation notation*
  - provides a concise, unambiguous description of data element encoding
  - e.g.,

```
type ObjectType UnsignedInteger (32)

type MunitionQuantity sequence {
    munition ObjectType,
    quantity Float (32)
}
```

Object type numbering scheme

- object type codes describe varieties of vehicles, ammunition, fuel, repair parts...
- object type codes are arranged in a hierarchy
- the hierarchy, once defined, remains valid as new object types are added
- software can understand something about an object based on its place in the hierarchy
- new types of objects can be introduced without disrupting existing software

Elements of communication compatibility

- scope of simulation: what phenomena are simulated
- architecture: what is computed where
- messages and their information content
- message encoding
- use of underlying network services
- ongoing simulation internet administration
  - e.g., assignment of simulator and object type numbers

Future work

- extensions for additional types of vehicles and simulators
- dead reckoning algorithms
  - using higher-order derivatives of location and rotation
  - blending in new appearance information
- missiles
  - transfer of missile simulation from firing simulator
  - homing on continuously designated targets
- dynamically changing terrain
- atmospheric effects, such as smoke and weather

The SIMNET Application Gateway

SIMNET WAN Peer Review

1 March 1990

Rolland M. Waters

BBN Systems and Technologies

Advanced Simulation Division

RWaters@bbn.com

(617) 873-2519

Uses of the SIMNET WAN

What	Timeliness	Number Interacting	Mode of Interaction	Type of Data
Vehicle Appearance	Real-Time	Many	Non-Reliable Sequential	Sequence small blks
Exercise Key Event Notification	Time Critical	Two	Reliable, Transaction	Small blocks
FM Voice Channels	Real-Time	One to Many	Non-Reliable, Shared	Sequence small blks
Key Event Logging	Not Time Critical	Many to One	Sequential Transactions	Sequence small blks
Exercise Checkpointing	Not Time Critical	Many to One One to One	Reliable, Transaction	Medium blocks
Conferencing/ AAR Voice & Video	Real-Time	Many	Non-Reliable Sequential	Sequence small blks
Workstation-based Conferencing	Interactive	Many	Reliable Sequential	Sequence Med blks
File Transfer and Mail	Not Time Critical	Two, One to Many	Reliable, Bulk	Small to Large blks
Remote Login	Interactive	Two	Reliable, Sequential	Sequence bytes->blks

Requirements of SIMNET WAN

- Multicast
- High Bandwidth
  - 1 Mbps now, 10 Mbps in 1 to 2 years
- Real Time Delivery
  - 100 to 250 ms end-to-end delay
  - Minimal delay variability
- Low Packet Loss Rate
- Wide Area Coverage
  - CONUS, Europe, Global
  - Tens of large sites
  - Hundreds of smaller sites
- Security
- Robust and Reliable
  - Advanced network management and control
  - Guaranteed performance
  - Best effort datagram delivery
- Usable by a Variety of Applications
  - Interactive support
  - Network sharing

Exercise Traffic Patterns

- A 1000 vehicle exercise can produce:
  - 1.5 Mbps vehicle traffic @ 1000 pps
  - 300 kbps voice traffic @ 300 pps
- This is well above the T1 "knee"
- During an exercise the great majority of traffic (> 80%) is Vehicle Appearance PDUs
- This means that most traffic requires real-time, multicast delivery
- Video conferencing will probably not be used during exercises.

THE SIMNET APPLICATION GATEWAY

Why a SIMNET Application Gateway?

- Computation/ communication tradeoffs are substantially different between the WAN and the LAN
- Need to reduce contention and queueing delays on WAN
- Must not increase vehicle processing load on individual simulators by adding too many additional vehicles from WAN
- New SIMNET services (e.g. digitized voice) need to be provided over WAN
- WAN must be kept transparent to simulators

FUNCTIONS OF THE SIMNET APPLICATION GATEWAY

- Compression of Vehicle Appearance PDUs
- Outbound filtering to reduce load on WAN
  - Includes RVA algorithms inappropriate for individual simulators
- Inbound filtering to reduce load on WAN
- Digitizing FM voice traffic
- Delay compensation
- Transparent interface to WAN

Vehicle Appearance PDU Compression

Vehicle Apperance PDU Contents:

- Static information	20%
- Slowly changing information	20%
- Dynamic information	60%

The static and slowly changing information is included in the PDU to make the protocol robust and fully distributed. The SIMNET Application Gateway can use inter-gateway protocols to update this information over the WAN so it doesn't have to be sent with each packet.

The dynamic information can be compressed by a factor of 3 through *floating-point intensive processing.*

Remembering that VAP traffic is more than 80% of the vehicle traffic:

=> The combination of these two techniques will result in a reduction of vehicle traffic of more than 64%.

=> This drops the typical traffic of a 1000 vehicle exercise to 800 kpbs, well below the T1 "knee".

Other Means of Reducing Traffic

- Outbound filtering
  - Don't send what the rest of the world doesn't want
- Inbound filtering
  - Don't fill your tail with traffic you don't want
- Better RVAs to reduce number of packets per vehicle
  - Higher order RVAs using acceleration and rotation in body coordinates
  - Use terrain for ground vehicle RVAs
  - Vehicle specific RVAs to account for turrets, etc
  - All of these are too expensive to implement in simulators
- Model static vehicles at remote sites
  - As RVAs get better, support vehicles will comprise a larger portion of vehicle traffic

Multicast And Real-Time Requirements

- Vehicle Appearance, Impact PDUs (90% of vehicle traffic)
  - Need both multicast and real-time delivery
  - End-to-end delay of 100ms acceptable for most vehicles
  - Maximum acceptable not known at present
    - Limited studies on tanks indicate 400ms acceptable
    - Air combat has tight limits
  - EW models may have lower acceptable delay

- Voice Traffic
  - Multicast
  - Almost as real-time as vehicle appearance

Inter-SIMNET Application Gateway traffic:

- Static information updates:
  - Time Critical
  - Reliable
  - Multicast

- Information for filtering
  - Reliable
  - Time sensitive
  - Multicast

- Advanced RVA Information
  - Real Time
  - Multicast
  - Needs modified form of reliability

Future WAN Requirements

1. Support Very Large Scale Exercises: Corps and Echelon Above Corps (10,000 to 30,000 vehicles)

   Need intelligent filtering/routing of multicast traffic in WAN: only send packets to sites which could be affected by them

2. Support Multimedia Conferencing for exercise prebrief, conferencing over topo maps

   Need shared workspace: sketch overlays on stored maps, converse about sketches

3. Support video teleconferencing for after-action review, informal exchange of ideas and plans among participants

   This has been found to be a major benefit to participants at NATO wargames

4. Support interconnect with real command and control consoles

   Convert to Simnet protocols to permit real equipment to be a full player in the Simnet battlefield

SIMNET Wide Area Network (WAN)

Claudio Topolcic

BBN Systems and Technologies Corporation

Interconnection to SIMNET WAN

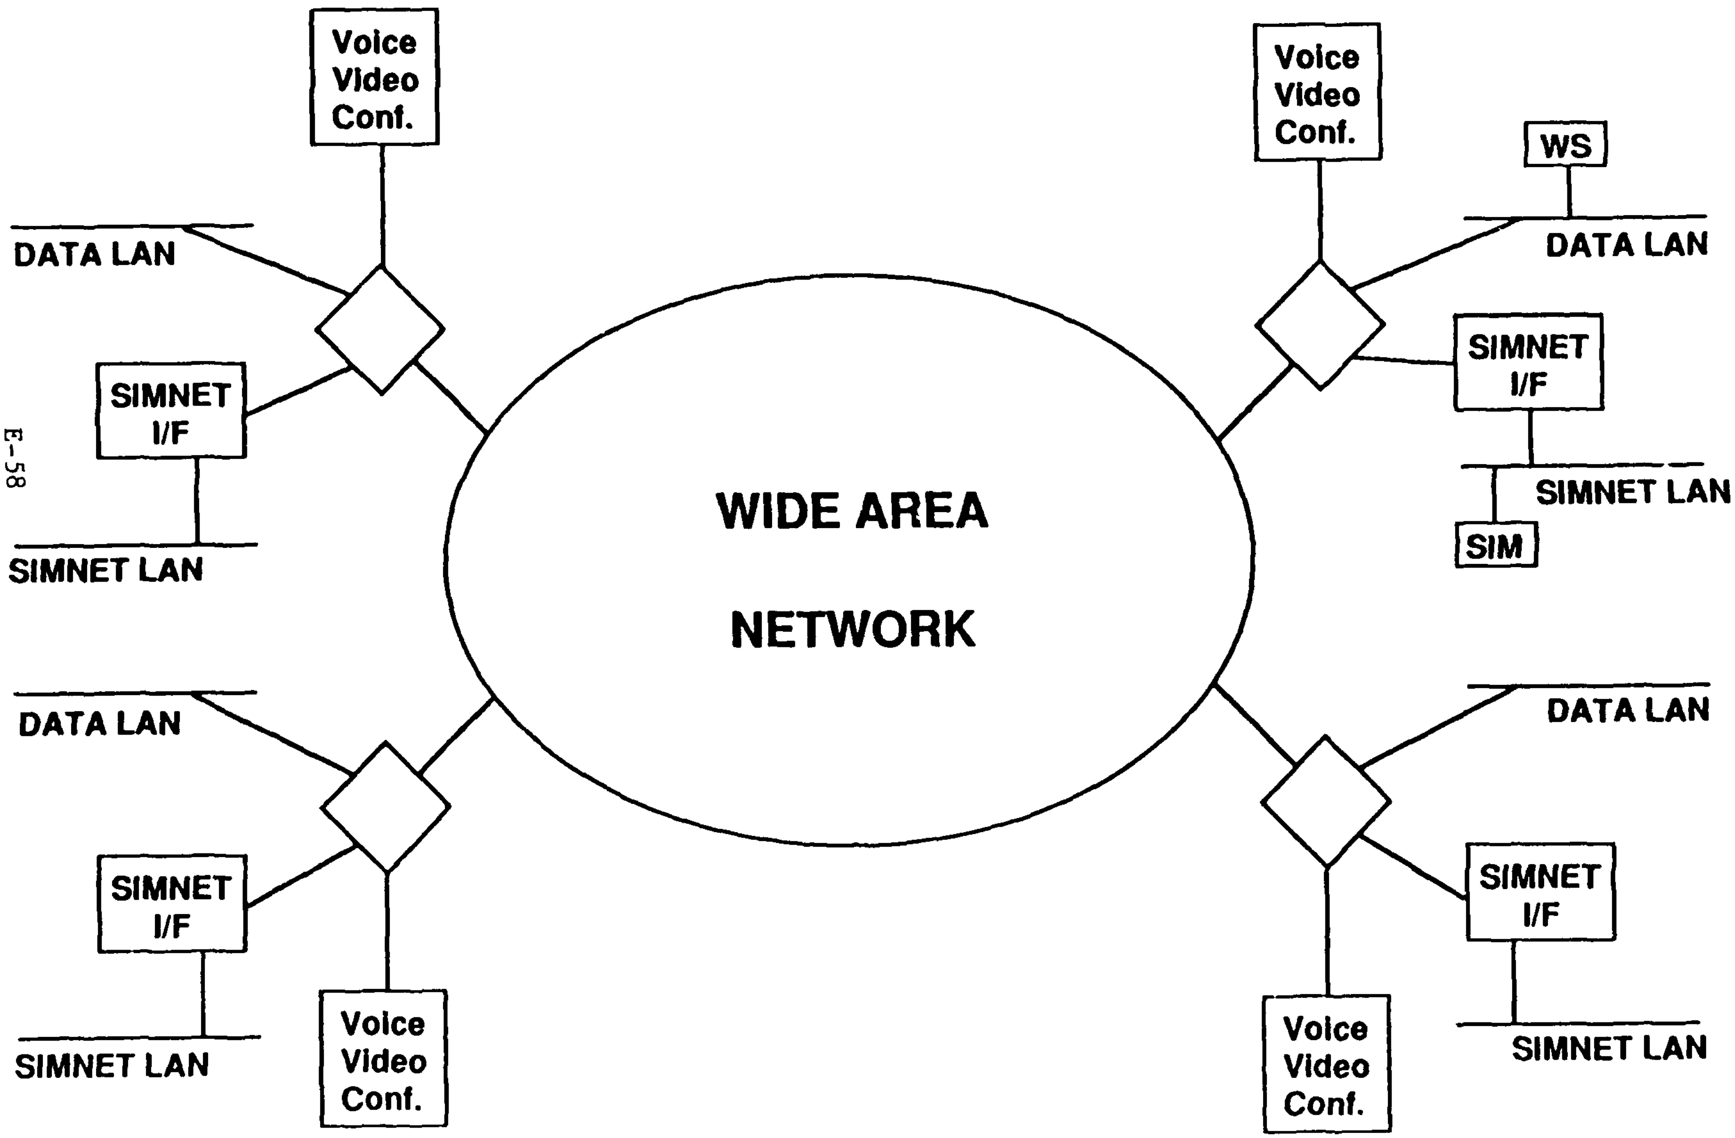

Current OSI Protocol Support for Real-Time Multicast Applications

- OSI Connection-Oriented Network Service (CONS/X.25)
  - No multicast
  - Window flow control adds delay
  - Retransmission for reliability adds delay
  - Does not provide performance guarantee

- OSI Connection-Less Network Service (CLNS)
  - No multicast
  - Requires non-standard transport protocol
  - Does not provide performance guarantee

- Future OSI Protocol May Support:
  - Multicast
  - Real-time transmission
  - Performance guarantee
  - Such a protocol would be ideal

SIMNET WAN Built of Internet Routers

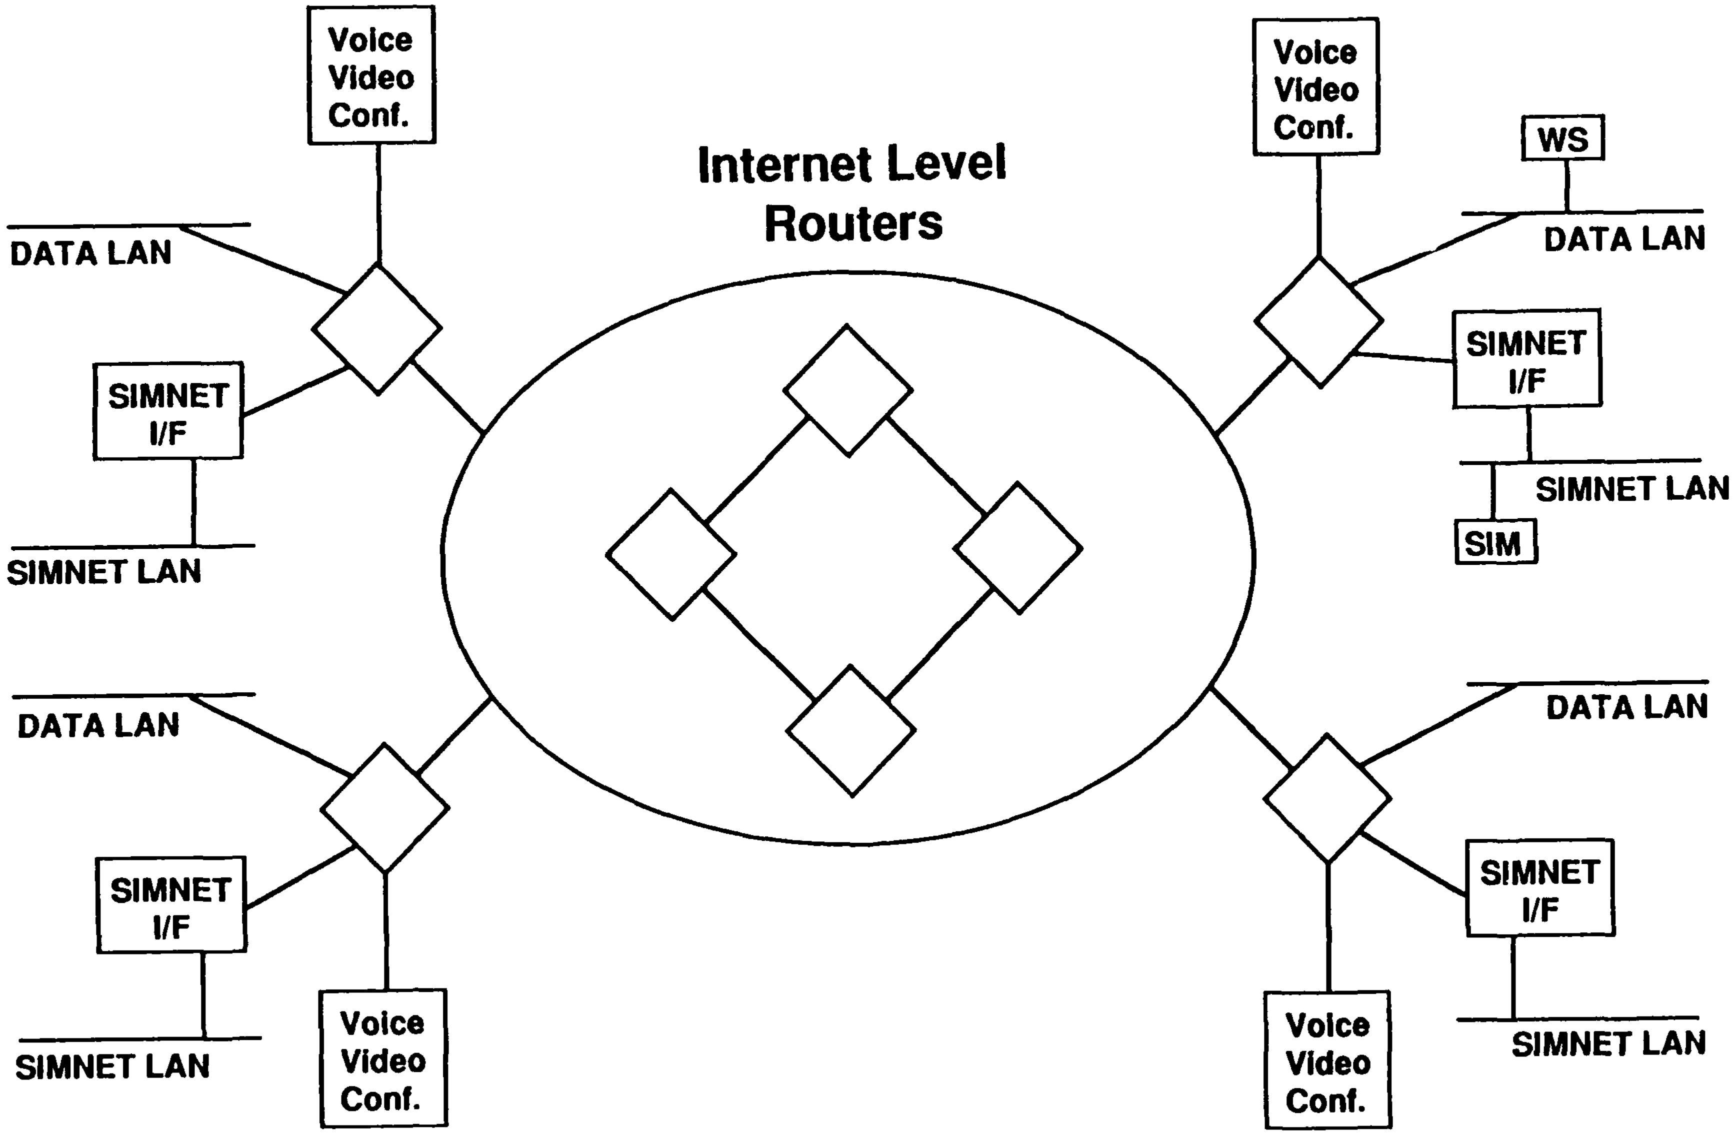

DoD Protocol Support for Real-Time Multicast Applications

- DoD Internet Protocol (IP) [Required Standard]
  - Supports multicast [Proposed]
  - Standard connection-less transport protocol available
  - Does not provide performance guarantees
  - Does not provide call blocking

- DoD Internet Stream Protocol (ST) [Experimental]
  - Supports multicast
  - Provides performance guarantees
  - Provides call blocking

- Other Evolving Internet Protocols [Under Development]
  - Zhang's Flow Protocol
  - Jacobsen's enhancements to IP
  - Follow-on to ST
  - May provide more uniform support for datagram and stream traffic
  - May be available in 2 to 3 years

Datagram vs. Stream Based Routing

Datagram

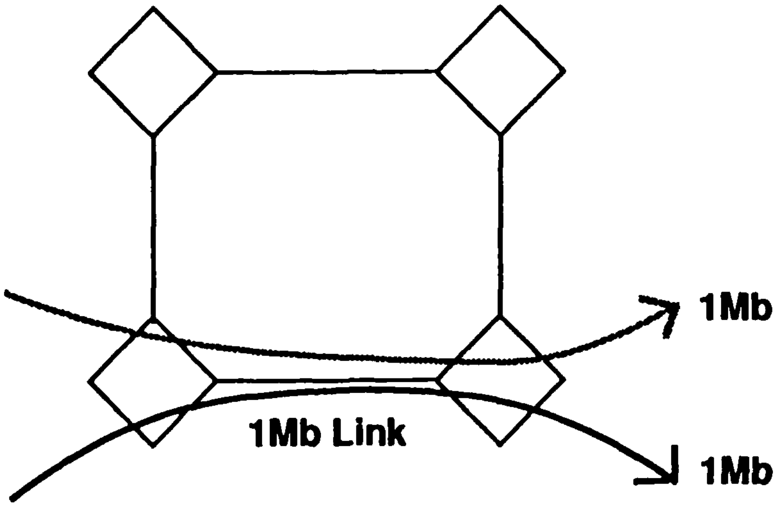

Stream or Flow

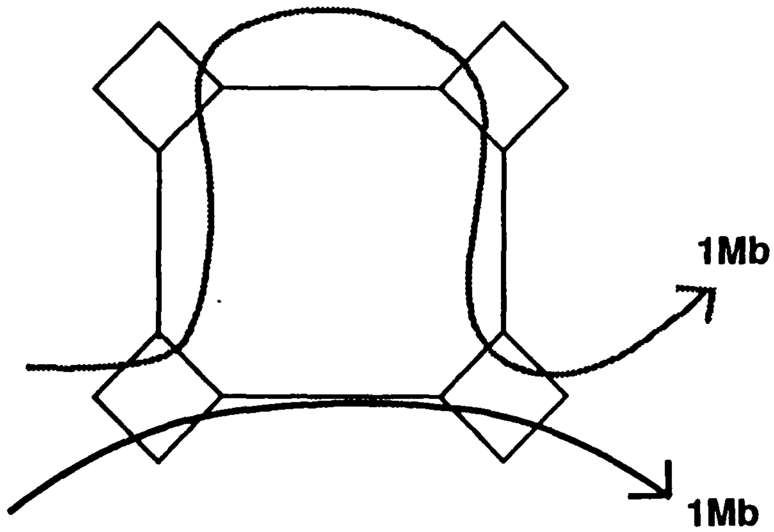

- Streams or Flows provide better support for performance guarantees
- Datagram routing is connectivity and hop-count based
- Datagram service does not take bandwidth requirements or bandwith availability into account
- Datagram routers cannot differentiate between packets in separate flows
- Streams or Flows accept requirements from application and routes based on requirements and availability
- Streams or Flows identify the stream to which a packet belongs

SIMNET WAN Built of IP/ST Routers

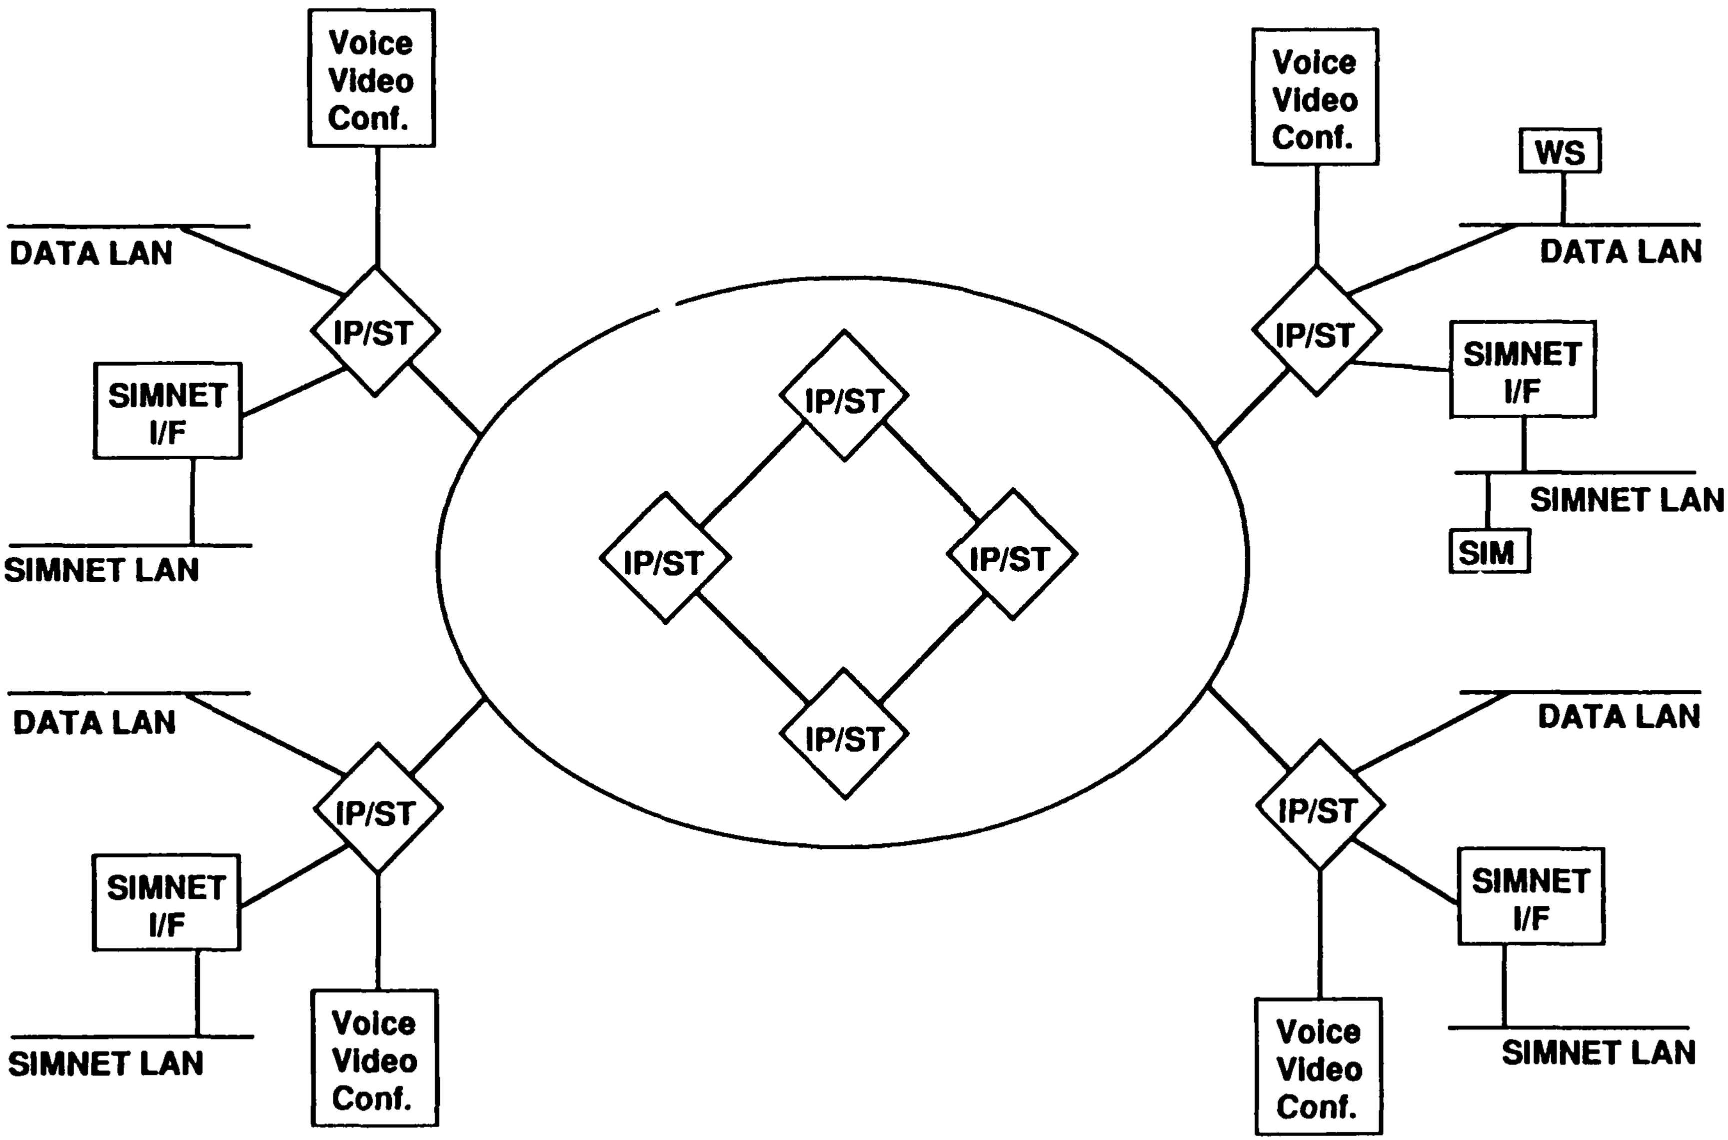

IP/ST Routers

- Advantages
  - Low risk
    - IP is a widely supported standard
    - ST is intended to support real-time multicast applications
    - An IP/ST Router currently exists
  - ST is the best support for next two or three years
  - ST is an internet experimental standard
- Disadvantages
  - ST is not yet a widely supported standard
  - ST may evolve or change in two to three years

Network Component in SIMNET WAN

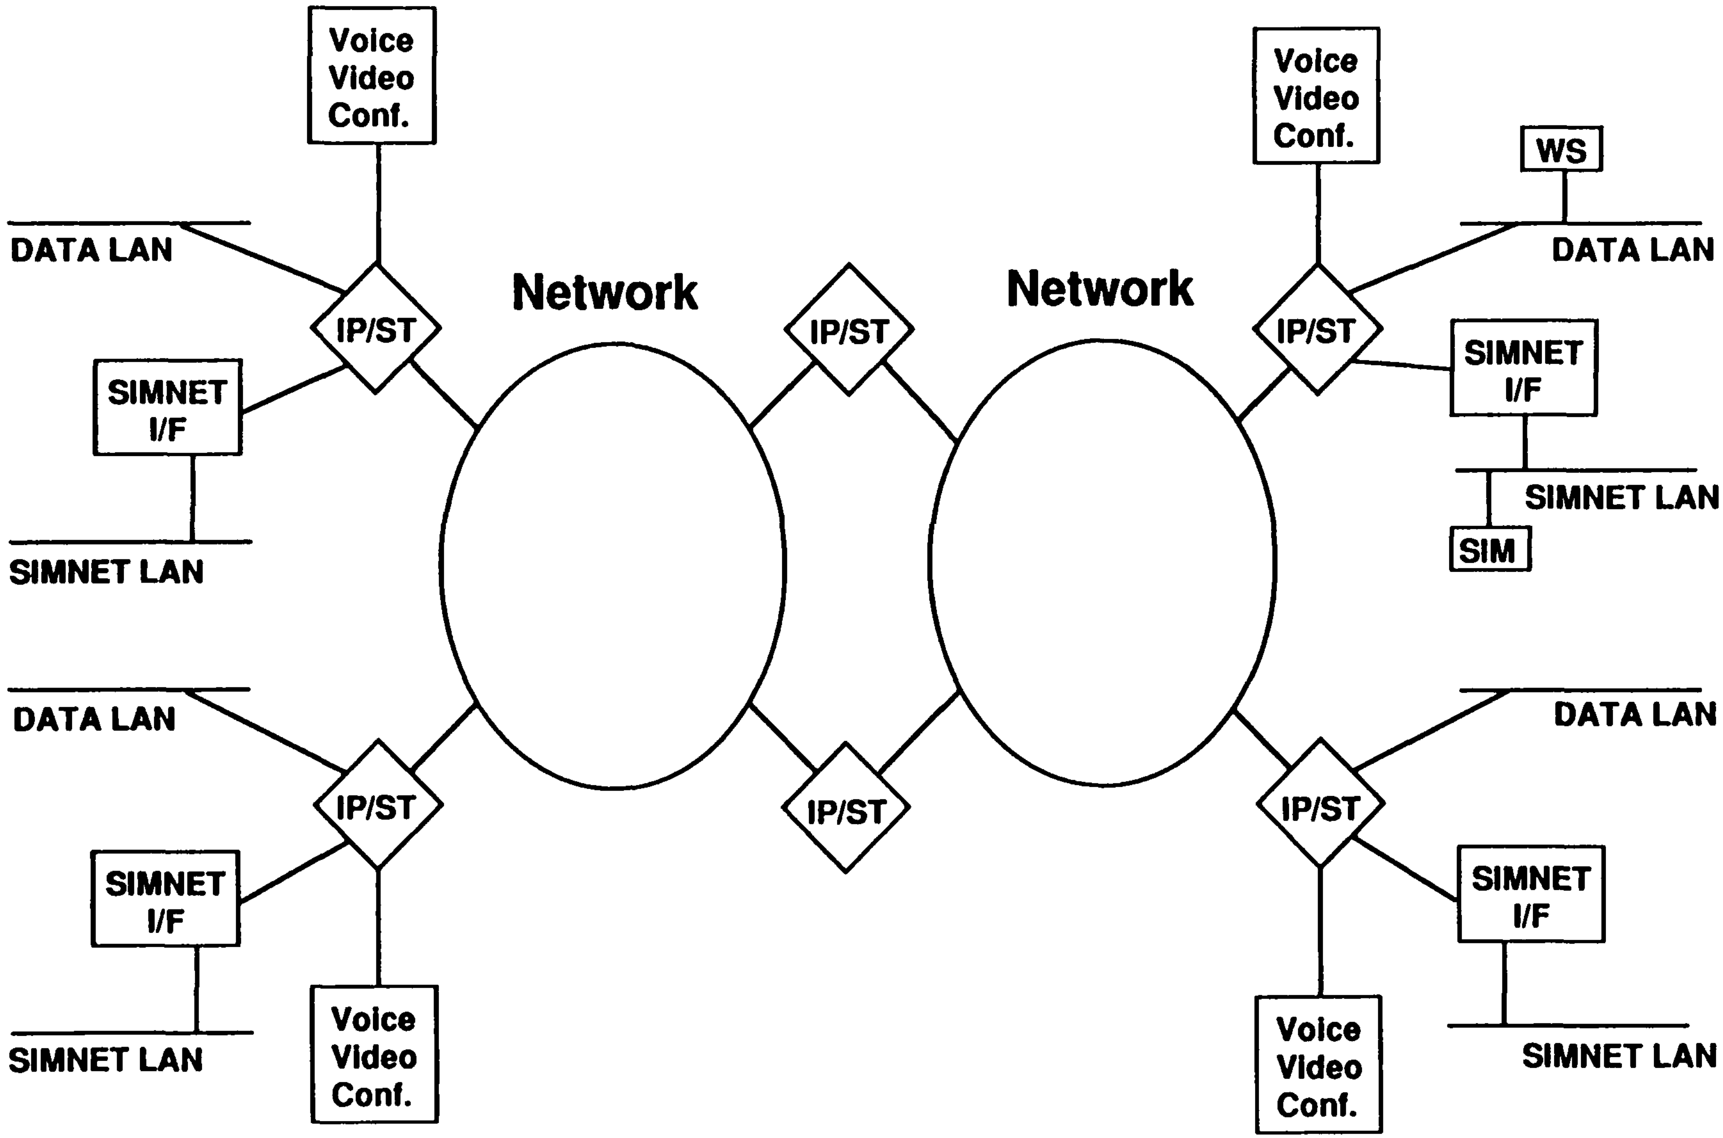

Network Component

- Advantages
  - Efficiency — packet forwarding can be faster within a network than at the internet level
  - Delay — packets can traverse large distances with minimal delay and overhead
- Disadvantage
  - Requires a separate system

BISDN Network Component

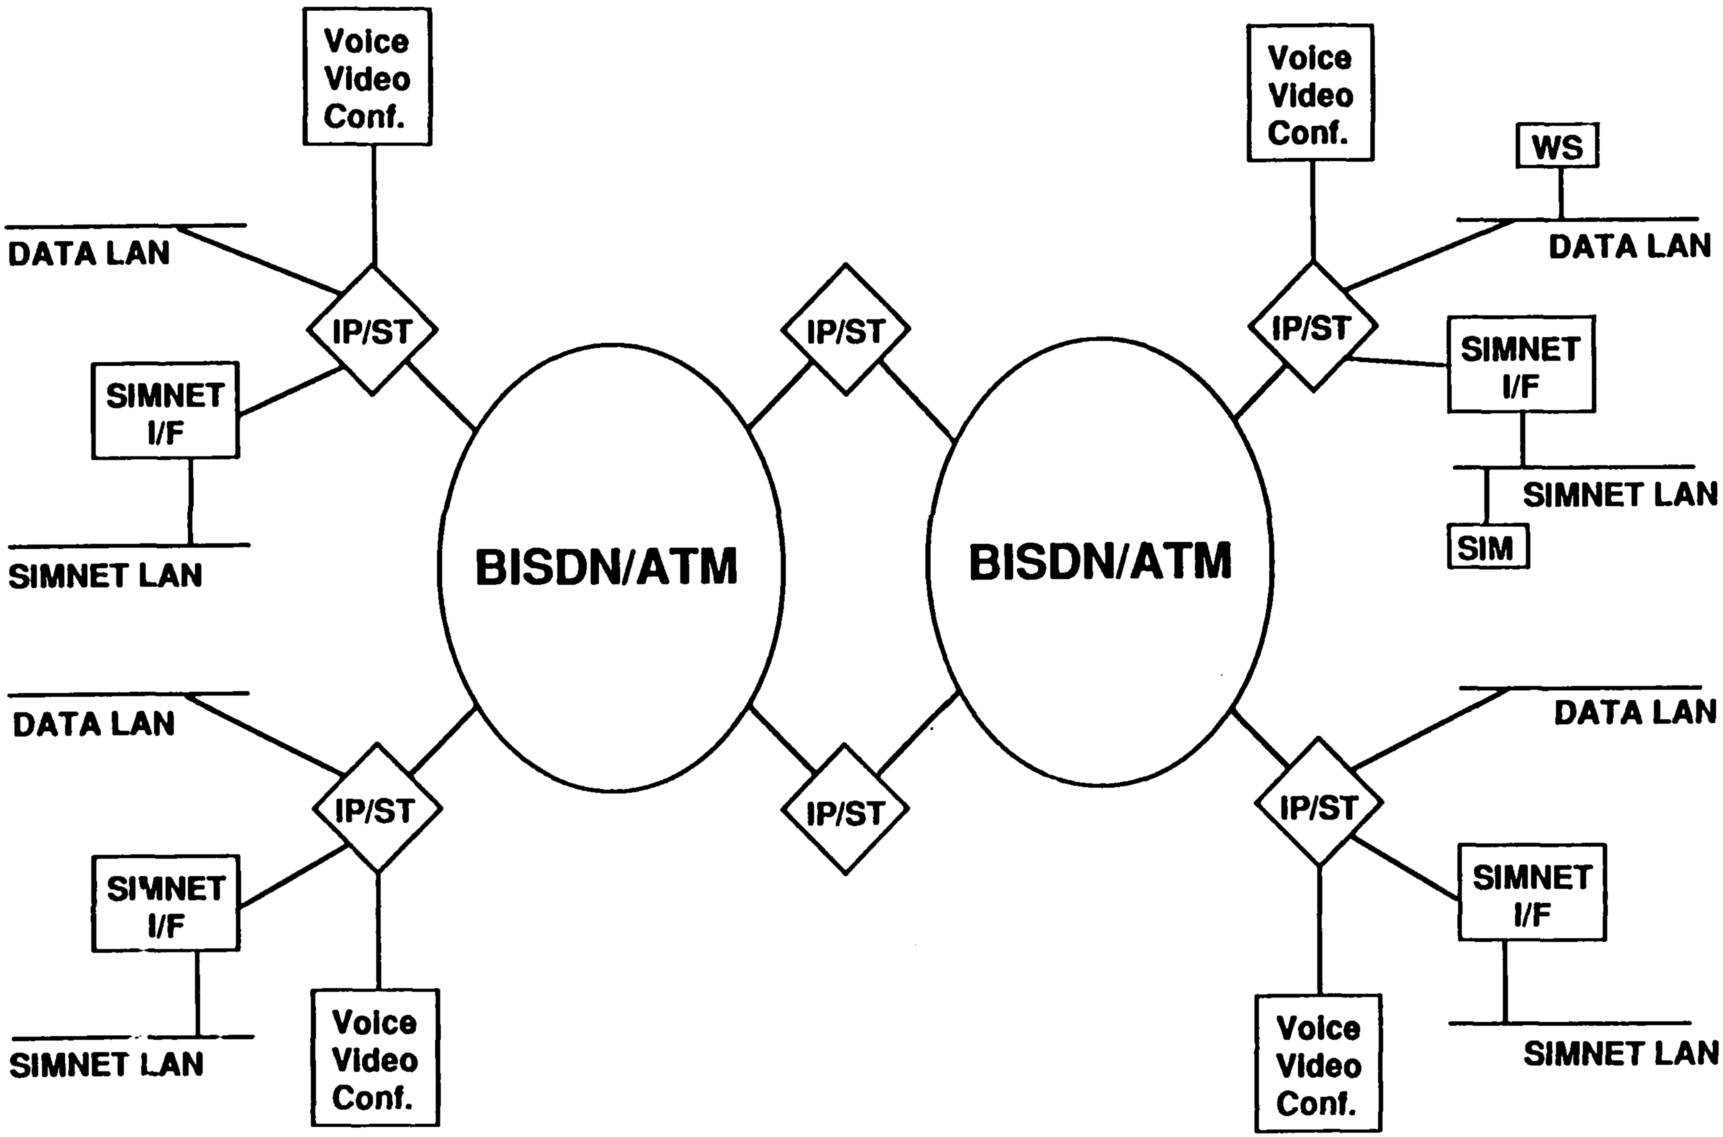

BISDN Network

- Advantages
  - Will be ubiquitous
  - Will have high performance
  - Promises to support real-time multicast

- Disadvantages
  - Real-time multicast support not expected for five to ten years
  - Specifics of control and performance have not yet been agreed to

Terrestrial Wideband Net Component of SIMNET WAN

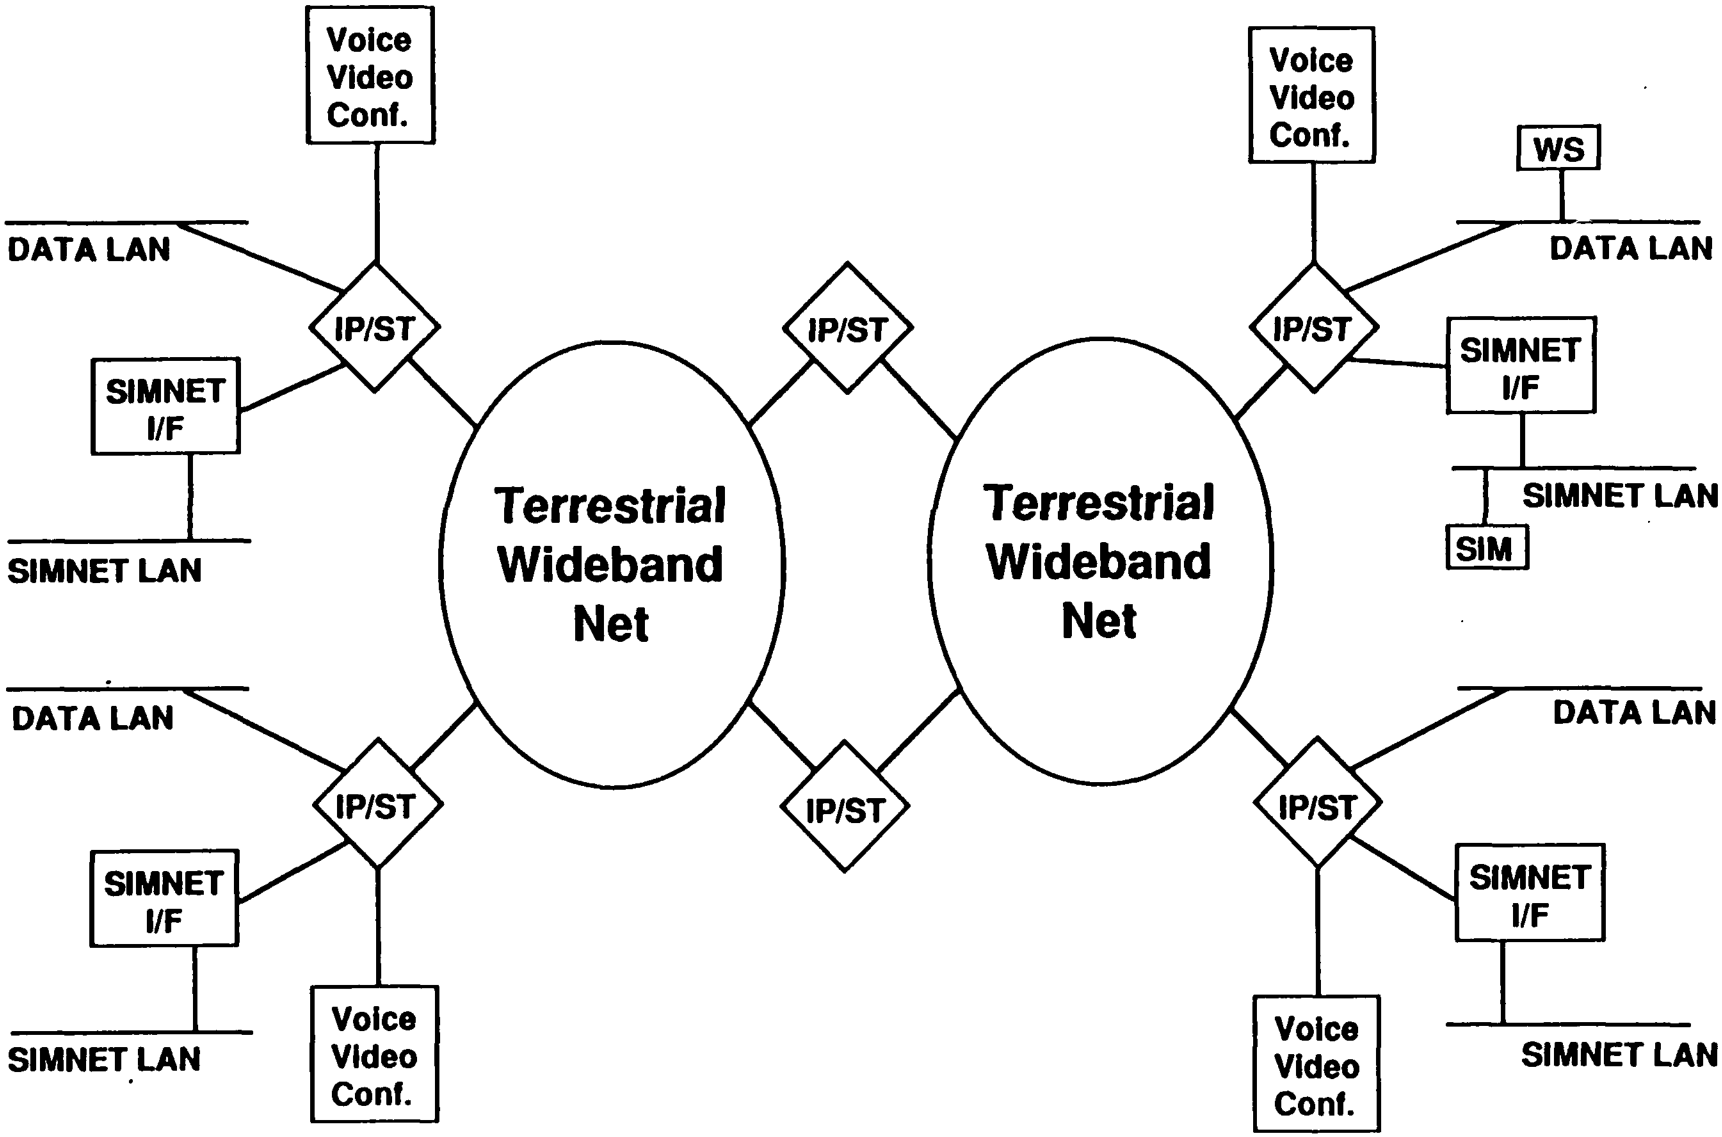

Terrestrial Wideband Net
(TWB Net)

- Advantages
  - Low risk
    - TWB is intended to support real-time multicast applications
    - TWB network currently exists
  - TWB technology provides lower forwarding delay than any internet router technology

- Disadvantages
  - Not yet a widely supported standard
  - Requires some further development to support a more complex topology

Proposed 2-3 year SIMNET WAN Architecture

- IP/ST routers connect SIMNET sites to WAN
- One or more Terrestrial Wideband Net provide wide area geographical coverage
- IP/ST routers interconnect the Terrestrial Wideband Networks
- IP/ST routers may also be directly interconnected
- Continue to work on protocol standarization

Proposed SIMNET WAN Architecture

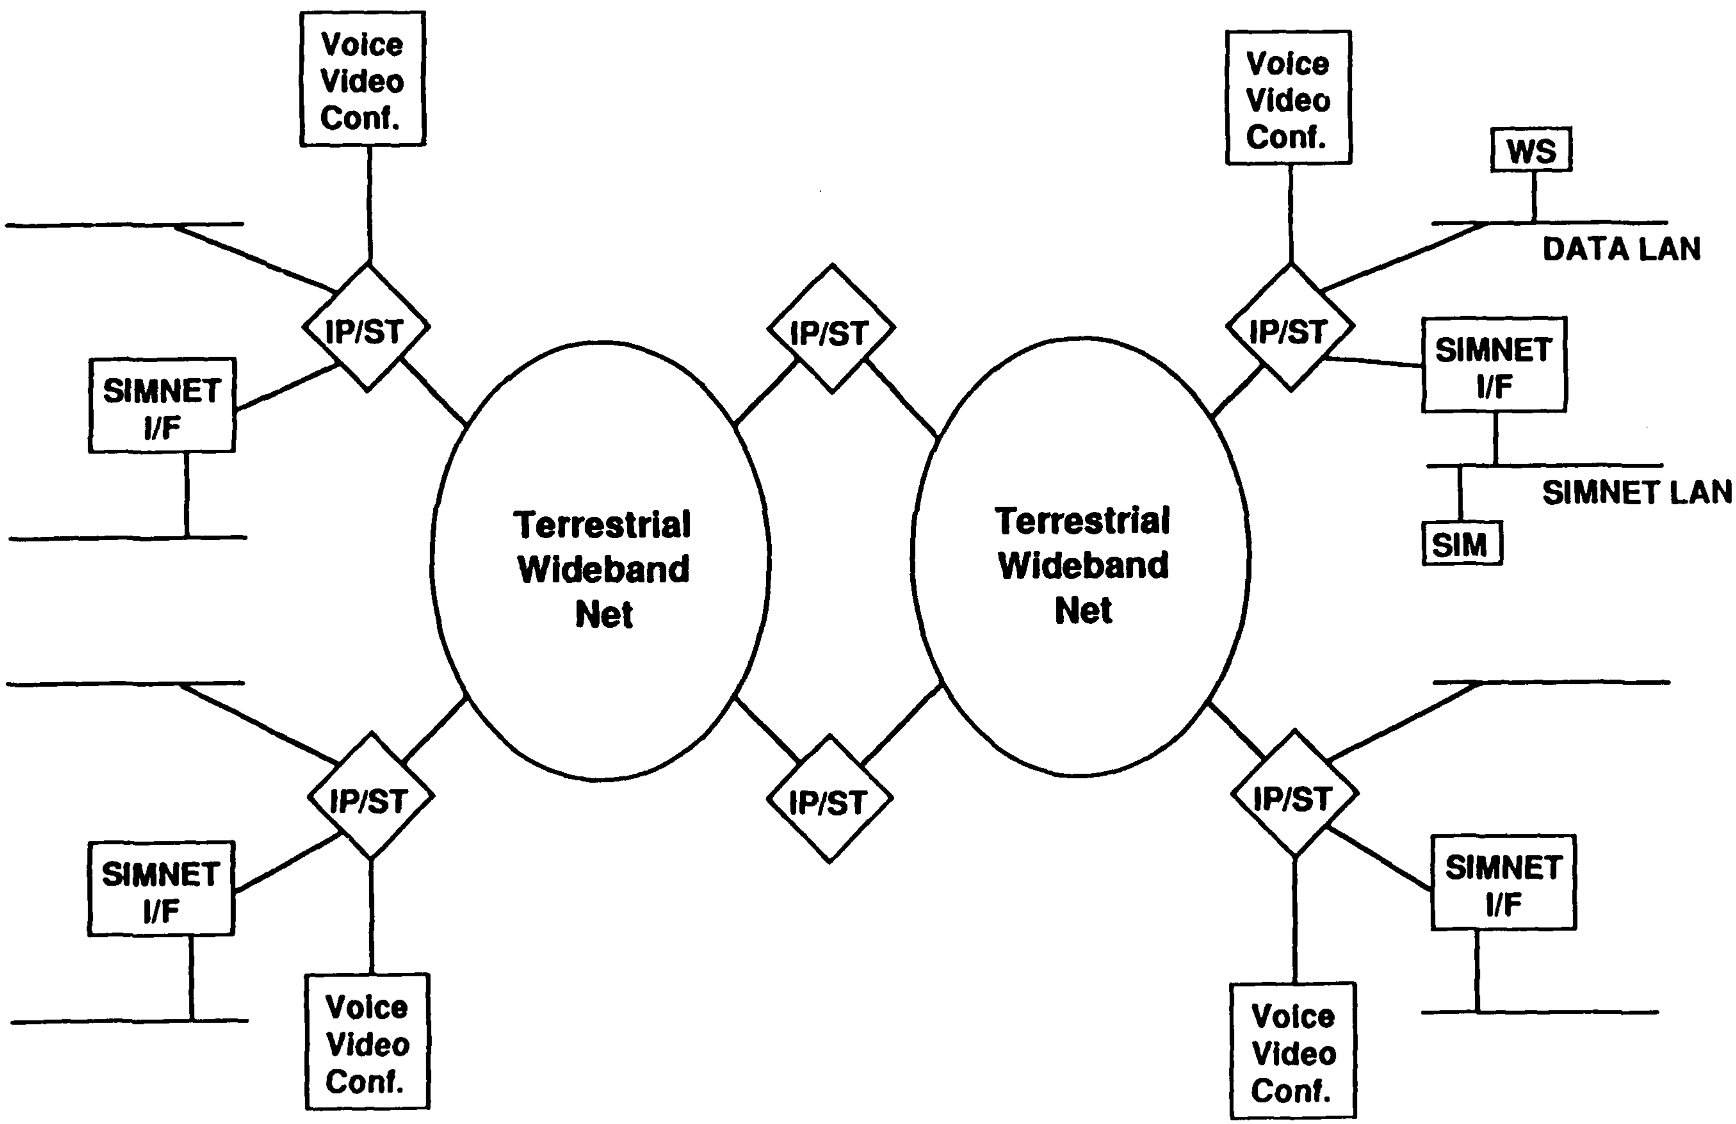

WAREX 3/90 Demonstration

- 800 vehicles distributed across 5 sites
- Vehicle types include:
  - Armored Vehicles
  - Helicopters
  - Fixed wing Aircraft
  - Support vehicles
- Other traffic:
  - packet voice - 10 channels @ 16Kb/s per channel

Terrestrial Wideband Network

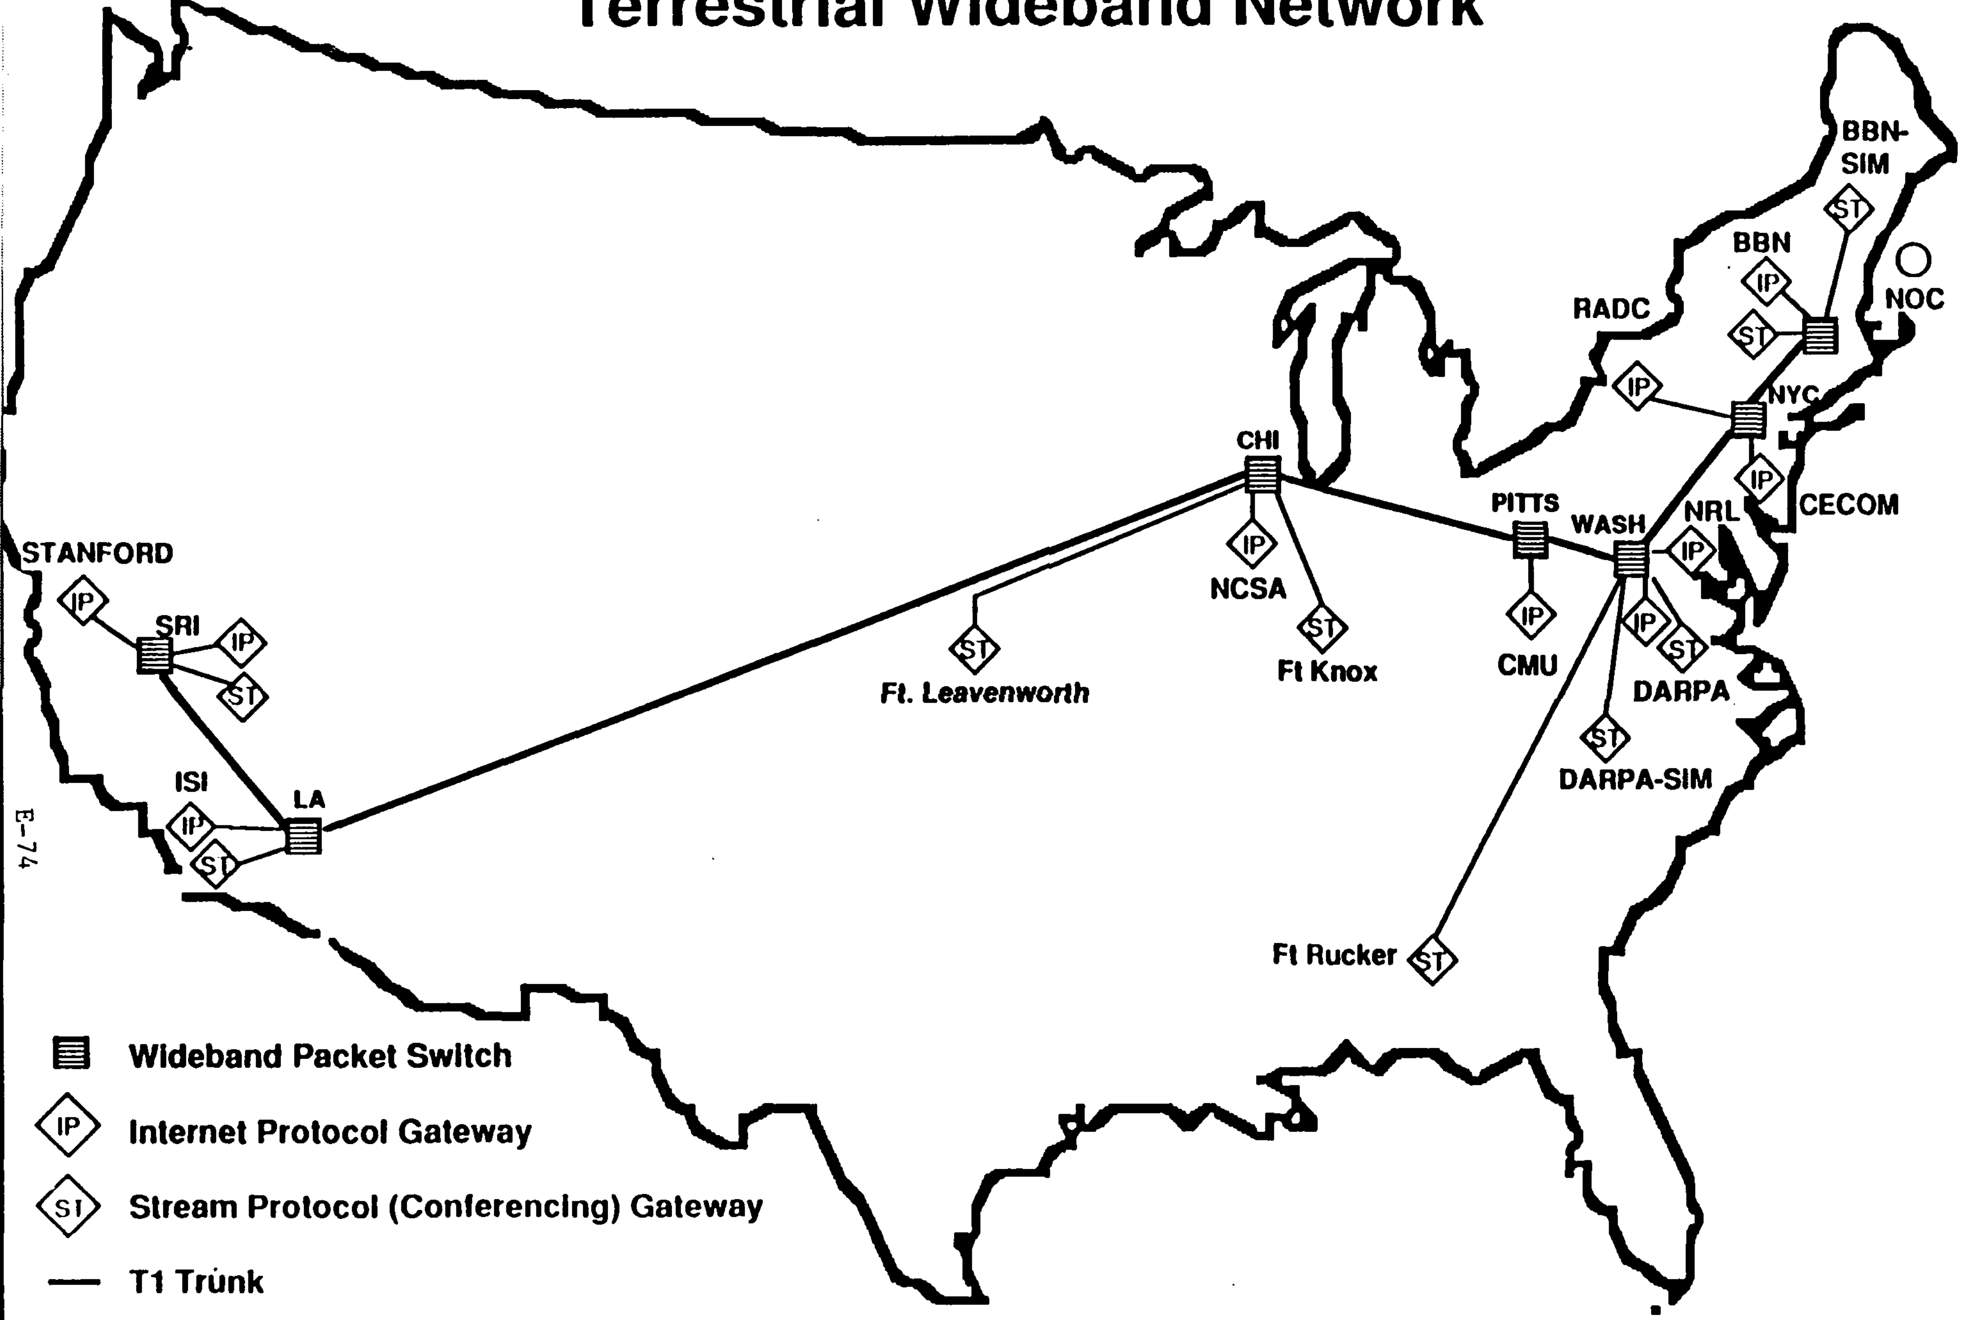

WAREX Site Detail

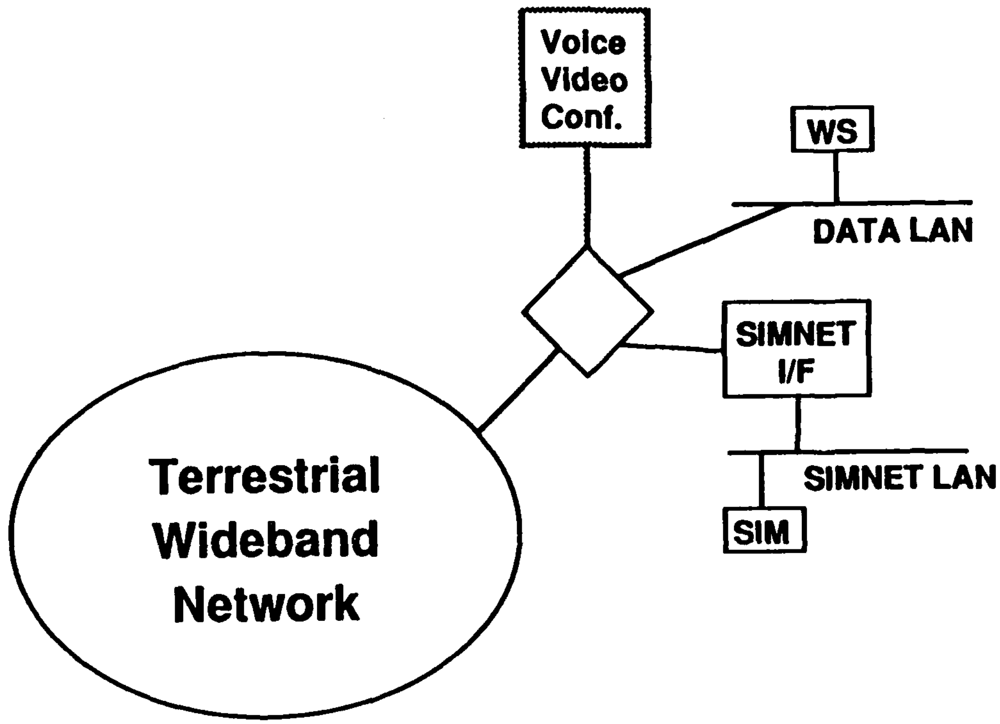

APPENDIX F
LONG-HAUL NETWORK BRIEFING PRESENTED BY MARTIN MARIETTA/SSDS

DSA

Distributed Simulators Architecture

Presentation for the Blue Ribbon Panel Review of Long-Haul Networking in Advanced Distributed Simulation Technology (ADST)

L. Michael Sabo

SSDS, Inc

2 MAR 90

AGENDA

- What is in the Governments Best Interest?
- Activities to Date
- The SIMNET Protocol Suite
- Why the SIMNET Protocols are Inadequate for the Future
- Why Standards and an Open Systems Approach?
- Overview of DSA
- Question, Comments, and Concerns

What is in the Governments Best Interest?

- The Ultimate Goal is a Migration to OSI
- The Government Does NOT Need a Custom Protocol Solution
    - Excessive Costs
    - Personnel Required to Tune the Protocols
    - Training on the Protocols
    - Porting Between Platforms
    - Software Maintenance

Activities to Date

- SIMNET Research Program Created as a Proof of Concept
- SIMNET Protocol Suite Created as Part of the Effort
- SIMNET Protocols Developed Without Scrutiny From Either the Internet or OSI Communities
- Protocols Stand as a *Single Vendor Implementation*
- First Workshop on Standards for Interoperability of Defense Simulators Held in August 1989
    - Told That SIMNET Works
    - We Were urged to Adopt it in its Totality

Activities to Date

- Second Workshop on Standards for Interoperability of Defense Simulators Held in January 1990
  - I Presented my Concerns to the Members
  - I was Asked to Prepare a Position Paper Regarding my Ideas
- I Prepared Four Position Papers Which Document DSA
- Panel Review for Long-Haul
  - Tall Pole in the Tent
  - Three Biggest Issues With Long-Haul
    - NATO, NATO, NATO
  - The Fundamental Problem is the SIMNET Protocol Suite

So Why was I Requested to Address This Panel?

- My Four Position Papers Describing DSA Became a Battle Cry
- I Have Received Encouraging Comments From Government, Industry, and Academia
- Distributed Simulation Technology is Moving Away From the Labs and into the Hands of the Military User

The SIMNET Protocol Suite

- The SIMNET Protocols are Much More Than a Simple Application Layer
- BBN Openly Admits That the "Association Protocol" Contains a Session and Transport Layer
- The Multicasting Capabilities are Rooted in Ethernet
- Few Operational Considerations Have Been Integrated
    - Network Management?
    - Security?
- Why is Long-Haul SIMNET a Problem?

The SIMNET Protocol Suite

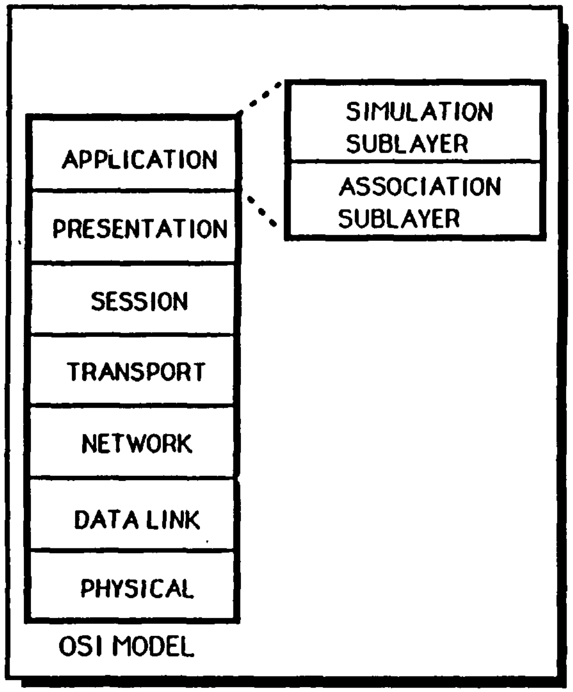

SIMNET ARCHITECTURE AS DOCUMENTED IN BBN REPORT NO. 7102

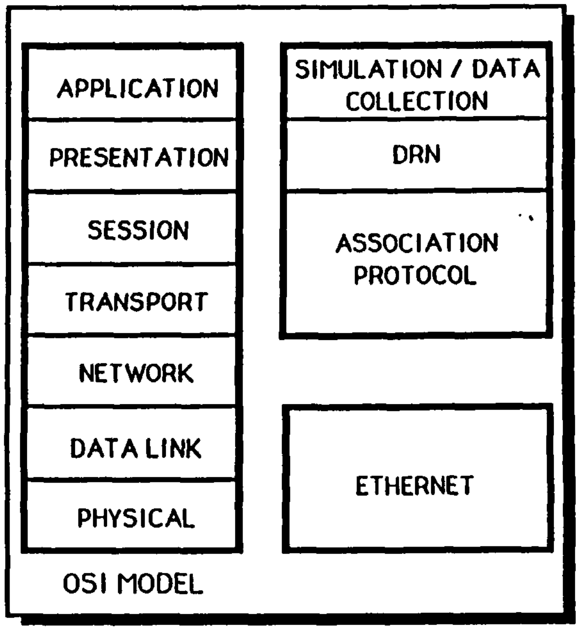

THE ACTUAL SIMNET ARCHITECTURE

DISTRIBUTED SIMULATORS ARCHITECTURE (DSA)

Why the SIMNET Protocols are Inadequate for the Future

- The Number of Vehicles on the Network is Expected to Grow Dramatically
  - No Routing Capability
  - No Provisions for Security
  - No Provision for Network Management
  - Inability to Interface With NATO
- SIMNET was designed as a Lab Experiment with Little or No attention Given to the Operational Environment
- Long-Haul Solution Must use Different Protocols than LAN Solution

Why Standards and an Open Systems Approach?

- Identical Architecture Across LANs, MANs, and WANs
- Draw on Large Technical Base for Solutions Such as Internet and OSI Communities
- Off-the-Shelf Equipment and Protocols Means Less $$$$$
- Modular or Layered Approach Facilitates Technology Insertion
- Our Problem is NOT Unique and we can Achieve Performance as well as Interoperability for the Defense Services

Overview of DSA

- DSA is Based on the OSI Seven Layer Reference Model
- Truly Transportable Application Layer Derived From SIMNET Simulation Protocol
- Lower Layers Will Use Off-The-Shelf Protocols
- DSA Will be Administered Through the Working Group Concept

Interactive Simulation Protocol (ISP)

- ISP Will be Implemented as an Open Application Layer Protocol
- ISP Will use the Presentation Services of ASN.1 Rather than DRN
- ISP will be Administered Through the ISP Working Group
- ISP will use the Simulation Protocol Portion of SIMNET as its Baseline
- Additional PDUs, as a Result of Interoperability Testing, may be Required
- The Public Domain Version of the Software will be Available on the Internet

HOW DO WE GET TO ISP?

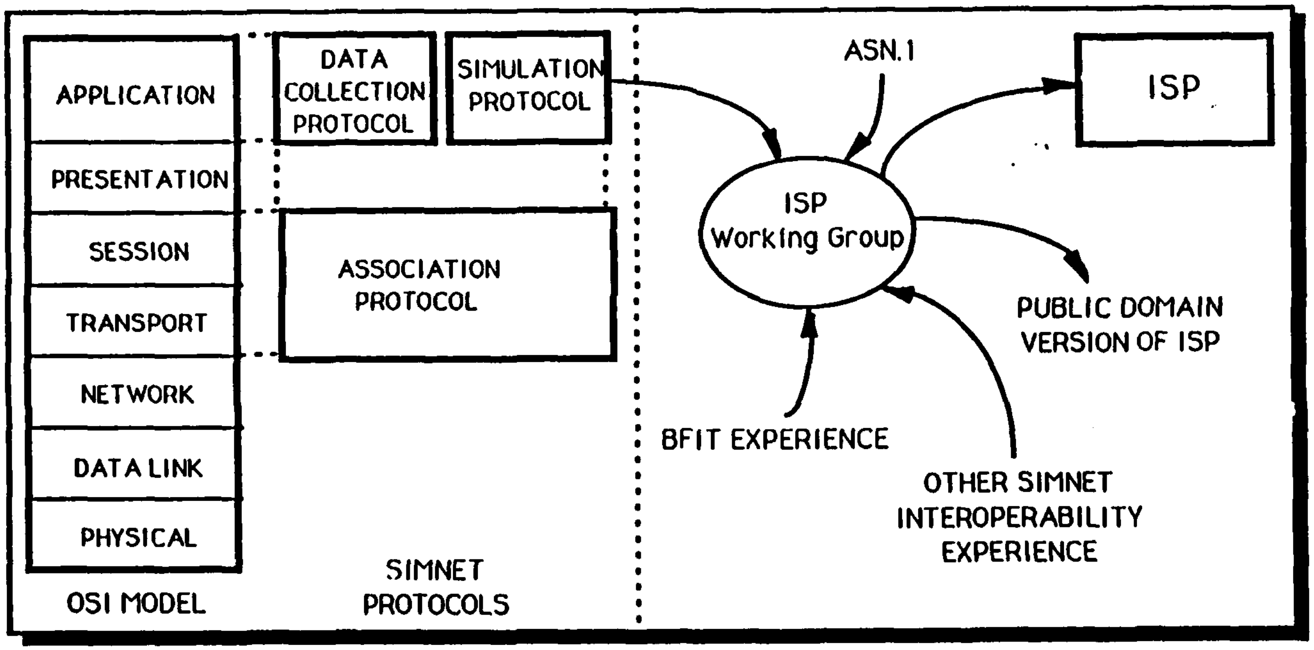

Figure 2 - ISP Development Methodology

DSA Transport Layer

- Transport Layer Provides End-to-End Logical Connection Between Simulators
- Transport Layer Insulates Upper Layer Protocols From the Network Particulars
- To be Considered a Viable Candidate the DSA Transport Layer Must Provide
    - Multicasting Capabilities
    - Datagram Service
- Three Potential Candidates Have Been Identified
    - Xpress Transfer Protocol (XTP)
    - Versatile Message Transaction Protocol (VMTP)
    - Multicast Internet Protocol (IP)
- XTP and VMTP Have Both Been Formally Presented to ANSI

Network Management

- Based on the Simple Network Management Protocol (SNMP)
- Managed Objects are Contained in a Logical Data Store Called the Management Information Base (MIB)
- Instrumentation Within the Network Node Collects Information for the MIB
- An SNMP Agent Operates Within a Network Node
- An SNMP Manager Operates Within a Network Management Station
- Agent Transfers MIB Data at the request of the Manager

Data Collection

- SNMP can easily be used for this function
- Private Enterprise MIB Established for This Function
- SNMP-Trap PDUs From Simulator to SIMNET MCC
- The Extensibility of the MIB allows Experimentation
- Why Build and Maintain Another Protocol When one With More Functionality is Available off the Shelf?

Summary

- SIMNET Protocols work but are not Suited to an Operational Environment
- DSA is Founded on Proven Principle of How a communications Architecture is Designed and Administered
- DSA is an Order of Magnitude More Open Than the SIMNET Protocols
- DSA can Provide the Required Performance to Support Distributed Simulation
- DSA will Provide a Bridge to the Ultimate Goal of OSI